Texture Recognition Using Robust Markovian Features

Pavel Vácha and Michal Haindl

Institute of Information Theory and Automation of the ASCR,
Pod Vodarenskou Vezi 4, 182 08 Prague, Czech Republic
{vacha,haindl}@utia.cas.cz

Abstract. We provide a thorough experimental evaluation of several state-of-the-art textural features on four representative and extensive image databases. Each of the experimental textural databases ALOT, Bonn BTF, UEA Uncalibrated, and KTH-TIPS2 aims at specific part of realistic acquisition conditions of surface materials represented as multispectral textures. The extensive experimental evaluation proves the outstanding reliable and robust performance of efficient Markovian textural features analytically derived from a wide-sense Markov random field causal model. These features systematically outperform leading Gabor, Opponent Gabor, LBP, and LBP-HF alternatives. Moreover, they even allow successful recognition of arbitrary illuminated samples using a single training image per material. Our features are successfully applied also for the recent most advanced textural representation in the form of 7-dimensional Bidirectional Texture Function (BTF).

Keywords: texture recognition, illumination invariance, Markov random fields, Bidirectional Texture Function, textural databases.

1 Introduction

Recognition of natural surface materials from their optical measurements represented as image textures, together with image (texture) segmentation, are the inherent part of plethora of computer vision algorithms, which are exploited in numerous real world applications such as visual scene analysis, image retrieval, medical images segmentation, image compression, etc. The key issue in solving real applications is robustness of employed methods, since images are usually captured in real non-laboratory environment, where acquisition conditions such as illumination, camera position, or noise cannot be controlled.

In this paper we focus on robustness of textural features to variations of illuminations conditions, such as spectrum, direction, and inhomogeneity. Illustrative examples of such appearance variations are displayed in Figs. 3, 4, and 5. Possible theoretical approach to robust recognition is learning from images captured under a full variety of possible illuminations for each material class [20,16], but it is obviously impractical, expensive to acquire and compute, or even impossible, if all needed measurements are not available. Alternatively, a kind of normalisation can be applied, e.g. cast shadow removal [7] or [8], which, unfortunately,

E. Salerno, A.E. Çetin, and O. Salvetti (Eds.): MUSCLE 2011, LNCS 7252, pp. 126–137, 2012.

completely wipes out rough texture structures with all their valuable discriminative information. Finally, the last and widely used approach is to construct corresponding invariants, which are features that do not change under specific variations of circumstances. However, it is necessary to keep in mind that an overdone invariance to broad range of sensing conditions inevitably reduces discriminability of features.

One of popular textural features are Local Binary Patterns [14] (LBP), which are invariant to any monotonic changes of pixel values, but they are very sensitive to noise [18] and illumination direction [19]. The LBP-HF extension [1] studies also relations between rotated patterns. Noise vulnerability was recently addressed by Weber Local Descriptor [4] (WLD). Texture similarity under different illumination direction [5] require the knowledge of illumination direction for all involved (trained as well as tested) textures. Finally, the MR8 texton representation [20] was extended to be colour and illumination invariant [2].

Multispectral textures can be described either jointly by multispectral textural features or separately by monospectral features on intensity image and colour features without spatial relations (histograms). The separate representation was advocated by [11], but we oppose this since a separate representation is not able to distinguish textures differing in position of pixels which have the same luminance. Obviously, the colour invariants computed from joint textural representation utilize the whole available information and they can create robust and compact texture description.

The contribution of this paper is a thorough evaluation of leading textural features under varying illumination spectrum, direction and slight variation of a camera location. We also test robustness to different acquisition devices, which is relevant especially for content-based image retrieval. These extensive tests of state of the art features were performed on four textural databases differing in variation of acquisition conditions and the results confirmed outstanding performance of Markovian textural features, preliminary tested in [19].

2 Markovian Textural Features

Our texture analysis is based on spatial and multimodal relations modelling by a wide-sense Markovian model. We employ a Causal Autoregressive Random (CAR) model, because it allows very efficient analytical estimation of its parameters. Subsequently, the estimated model parameters are transformed into illumination / colour invariants, which characterize the texture. These colour invariants encompass inter-spectral and spatial relations in the texture which are bounded to a selected contextual neighbourhood, see Fig. 1.

Let us assume that multispectral texture image is composed of C spectral planes (usually $C = 3$). $Y_r = [Y_{r,1}, \ldots, Y_{r,C}]^T$ is the multispectral pixel at location r, where the multiindex $r = [r_1, r_2]$ is composed of r_1 row and r_2 column index, respectively. The spectral planes are either modelled by 3-dimensional CAR model or mutually decorrelated by the Karhunen-Loeve transformation (Principal Component Analysis) and subsequently modelled using a set of C 2-dimensional CAR models.

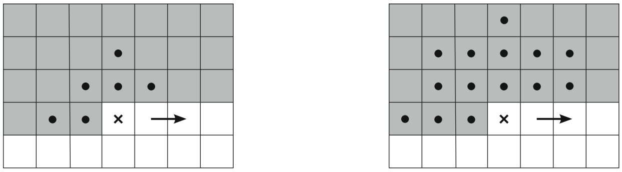

Fig. 1. Examples of contextual neighbourhood I_r. From the left, it is the unilateral hierarchical neighbourhood of third and sixth order. X marks the current pixel, the bullets are pixels in the neighbourhood, the arrow shows movement direction, and the grey area indicates acceptable neighbourhood pixels.

The CAR representation assumes that the multispectral texture pixel Y_r can be modelled as a linear combination of its neighbours:

$$Y_r = \gamma Z_r + \epsilon_r \ , \qquad Z_r = [Y_{r-s}^T : \forall s \in I_r]^T \tag{1}$$

where Z_r is the $C\eta \times 1$ data vector with multiindices r, s, t, $\gamma = [A_1, \ldots, A_\eta]$ is the $C \times C\,\eta$ unknown parameter matrix with square submatrices A_s. Some selected contextual causal or unilateral neighbour index shift set is denoted I_r and $\eta = \mathit{cardinality}(I_r)$, see Fig. 1. The white noise vector ϵ_r has normal density with zero mean and unknown full covariance matrix, same for each pixel.

The texture is analysed in a chosen direction, where multiindex t changes according to the movement on the image lattice. Given the known history of CAR process $Y^{(t-1)} = \{Y_{t-1}, Y_{t-2}, \ldots, Y_1, Z_t, Z_{t-1}, \ldots, Z_1\}$ the parameter estimation $\hat{\gamma}$ can be accomplished using fast and numerically robust statistics [9]:

$$\begin{aligned}
\hat{\gamma}_{t-1}^T &= V_{zz(t-1)}^{-1} V_{zy(t-1)} \ , \\
V_{t-1} &= \begin{pmatrix} \sum_{u=1}^{t-1} Y_u Y_u^T & \sum_{u=1}^{t-1} Y_u Z_u^T \\ \sum_{u=1}^{t-1} Z_u Y_u^T & \sum_{u=1}^{t-1} Z_u Z_u^T \end{pmatrix} + V_0 = \begin{pmatrix} V_{yy(t-1)} & V_{zy(t-1)}^T \\ V_{zy(t-1)} & V_{zz(t-1)} \end{pmatrix} \ , \\
\lambda_{t-1} &= V_{yy(t-1)} - V_{zy(t-1)}^T V_{zz(t-1)}^{-1} V_{zy(t-1)} \ ,
\end{aligned} \tag{2}$$

where the positive definite matrix V_0 represents prior knowledge.

In the case of 2D CAR models stacked into the model equation (1), the uncorrelated noise vector components ϵ_r are additionally assumed. Consequently, the image spectral planes have to be decorrelated before modelling and the parameter matrices A_s are diagonal (in contrast with full matrices for general 3D CAR model).

Colour Invariants

Colour invariants are computed from the CAR parameter estimates to make them independent on changes of illumination intensity and colours. More precisely, these invariants are invariant to any linear change of pixel value vectors $B\,Y_r$, where B is $C \times C$ regular transformation matrix. This is in accordance with reflectance models including specular reflections and even with the majority of

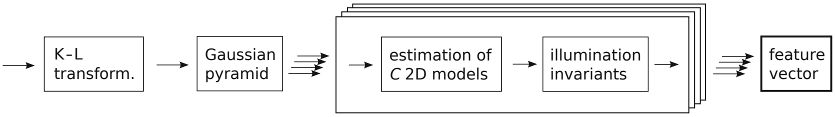

Fig. 2. Texture analysis algorithm using a set of 2D random field models

available BTFs [17], if illumination position remains unchanged. Additionally, 2D CAR models assume that the matrix B is diagonal. Moreover, our invariants are approximately invariant to infrequent changes of local illumination intensity and experiments show their robustness to variation of illumination direction. The following colour invariants were derived [18,17]:

1. trace: $\operatorname{tr} A_s,\ \forall s \in I_r$,
2. eigenvalues: $\nu_s = \operatorname{eigs}(A_s),\ \forall s \in I_r$,
3. α_1: $1 + Z_r^T V_{zz}^{-1} Z_r$,
4. α_2: $\sqrt{\sum_r \left(Y_r - \hat{\gamma} Z_r\right)^T \lambda^{-1} \left(Y_r - \hat{\gamma} Z_r\right)}$,
5. α_3: $\sqrt{\sum_r \left(Y_r - \mu\right)^T \lambda^{-1} \left(Y_r - \mu\right)}$, μ is the mean value of vector Y_r .

The model parameters $\hat{\gamma}, \lambda$ are estimated using formula (2), we omit subsctripts for simplicity. Feature vectors are formed from these illumination invariants, which are easily evaluated during the CAR parameters estimation process.

In the case of 2D models, no eigenvalues are computed because matrices A_s are diagonal, and the features are formed from the diagonals without their reordering:

2. diagonals: $\nu_s = \operatorname{diag} A_s\,, \quad \forall s \in I_r\,.$

Moreover, the invariants $\alpha_1 - \alpha_3$ are computed for each spectral plane separately.

Algorithm

The texture analysis algorithm starts with factorisation of texture image into K levels of the Gaussian down-sampled pyramid and subsequently each pyramid level is modelled by the CAR model. The pyramidal factorization is used, because it enables model to easily capture larger spatial relations. We usually use $K = 4$ levels of Gaussian down-sampled pyramid and the CAR models with the 6-th order semi-hierarchical neighbourhood (cardinality $\eta = 14$). If the image size is large enough (at least 400×400) it is possible to improve performance with the additional pyramid level ($K = 5$). Finally, the estimated parameters for all pyramid levels are transformed into the colour invariants and concatenated into a common feature vector. The algorithm scheme for 2D CAR-KL is depicted in Fig. 2, where "-KL" suffix denotes decorrelation by Karhunen-Loeve transformation.

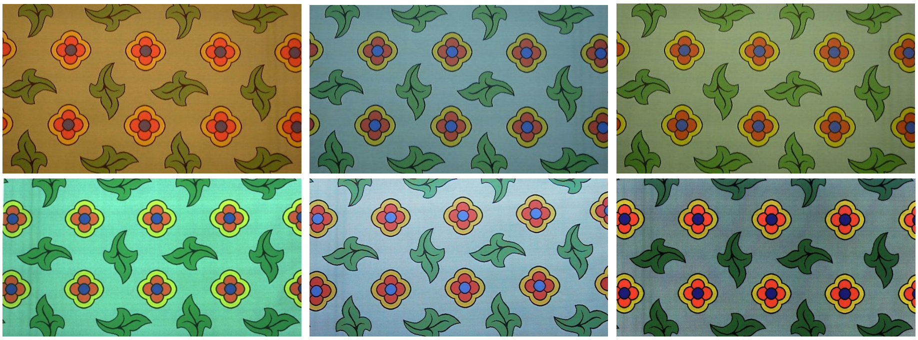

Fig. 3. Images from the UEA database, the upper row displays images under different illumination, while the bottom row shows images from different acquisition devices

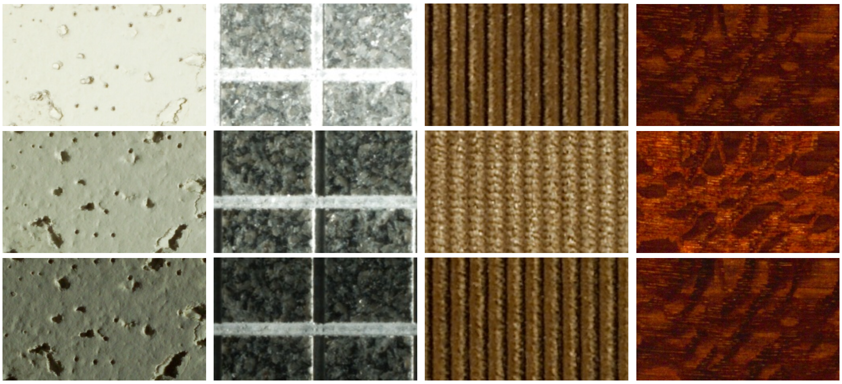

Fig. 4. Images from the Bonn BTF database, the first and second column with different illumination declination and the rest with various illumination azimuth

Fig. 5. Images from the ALOT database, each column shows images of the same material captured under varying illumination conditions

The dissimilarity between two feature vectors of two textures T, S is computed using fuzzy contrast [15] in its symmetrical form FC_3:

$$FC_3\,(T,S) = M - \left\{ \sum_{i=1}^{M} \min\left\{ \tau(f_i^{(T)}), \tau(f_i^{(S)}) \right\} - 3 \sum_{i=1}^{M} \left| \tau(f_i^{(T)}) - \tau(f_i^{(S)}) \right| \right\} ,$$

$$\tau(f_i) = \left(1 + \exp\left(- \frac{f_i - \mu(f_i)}{\sigma(f_i)} \right) \right)^{-1} ,$$

where M is the feature vector size and $\mu(f_i)$ and $\sigma(f_i)$ are average and standard deviation of the feature f_i computed over all database, respectively. The sigmoid function τ models the truth value of fuzzy predicate.

3 Experiments

We tested robustness of the CAR features in texture recognition on four different image data sets, each with different conditions. The first experiment is focused on recognition in variable illumination spectra and different acquisition devices, while the second experiment tests robustness to illumination direction changes. The next experiment utilises the largest recent collection of natural and artificial materials captured under various illumination conditions and the last test is classification into material categories. Summary of experiment setups is provided in Tab. 1.

Table 1. Parameters of experiments and comprised variations of recognition conditions

	Experiment			
texture database	UEA	Bonn BTF	ALOT	KTH
experiment conditions:				
illumination spectrum	+	−	+	−
illumination direction	−	+	+	−
viewpoint declination	−	−	+/−	−
acquisition device	+	−	−	−
experiment parameters:				
image size (bigger)	≈ 550	256	1536	200
number of classes	28	15,10	200, 250	11

The CAR features were compared with the most frequently used textural features as Gabor features [12], Opponent Gabor features [10], LBP [14], and LBP-HF [1]. These features demonstrated state of the art performance in the cited articles and all were tested with authors parameter settings. The grey level features such as Gabor features, LBP, and LBP-HF were computed either on grey level images or additionally for each spectral plane separately and concatenated, which is denoted with “RGB” suffix. For LBP features we tested variants

Table 2. Size of feature vectors

method	size	method	size
Gabor f.	144	2D CAR-KL	**260**
Opponent Gabor f.	**252**	3D CAR-KL	236
$LBP_{8,1+8,3}$,	512	2D CAR-KL (K=5)	325
$LBP_{8,1+8,3}$, RGB	**1536**	3D CAR-KL (K=5)	295
$LBP\text{-}HF_{8,1+16,2+24,3}$, RGB	1344	$LBP^{u2}_{16,2}$	243

$LBP_{8,1+8,3}$, $LBP^{u}_{16,2}$, and $LBP\text{-}HF_{8,1+16,2+24,3}$ reported by authors as the best in their experiments. Gabor features were additionally tested with and without separate normalisation of spectral planes (*Greyworld*), which is denoted with "norm." suffix. Size of feature vectors is summarised in Tab. 2. The following result figures display only the best performing features in each kind for each experiment.

3.1 University of East Anglia Uncalibrated Image Database

The first experiment was performed on UEA Uncalibrated Image Database[1] [6]. This dataset contains 28 textile designs, captured with 6 different devices (4 colour cameras and 2 colour scanners), and images for cameras were illuminated with 3 different illumination spectra, which sums up to 394 images in total. No calibration was performed and image resolution is about 550×450 (± 100). Examples of images are shown in Fig. 3.

In this experiment, training images per each material were randomly selected and the remaining images were classified using the Nearest Neighbour (1-NN) classifier, the results were averaged over 10^3 of random selections of training images. As it is displayed in Fig. 6, the alternative textural features were surpassed for all tested numbers of training images per material. It is quite surprising that LBP features had difficulties in this experiment, since they are invariant to any monotonic change of pixel values, while CAR features assume linear relation. UEA images are supposed to include even non-linear relations of images caused by different processing in acquisition devices. The poor performance of LBP features may be due to similarity of certain characteristics in UEA images that the LBP features are not able to distinguish or due to slight scale variation of images. The large images allowed to compute CAR features on $K = 5$ pyramid levels, the results for 2D CAR with $K = 4$ went from 56.6 to 85.3 for 1 to 6 training images, which still outperformed alternatives by a large margin.

3.2 Bonn BTF Database

The second experiment was performed on the University of Bonn BTF database [13], which consists of fifteen BTF colour measurements. Ten of those (corduroy, impalla, proposte, pulli, wallpaper, wool, ceiling, walk way, floor tile, pink tile)

[1] `http://www.uea.ac.uk/cmp/research/graphicsvisionspeech/colour/data-code/`

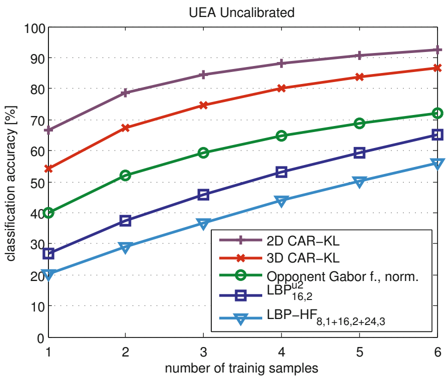

Fig. 6. Accuracy of classification on UEA images. The number of training images is changing from 1 to 6 per material.

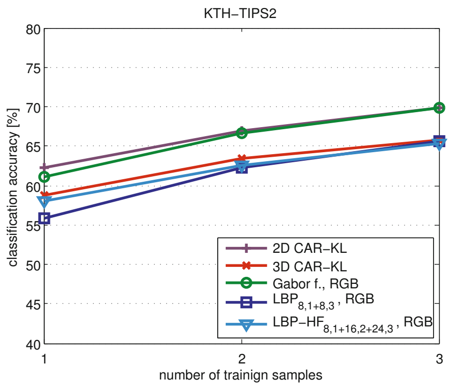

Fig. 7. Accuracy of classification on KTH_TIPS2 database. The number of training material samples is changing from 1 to 3 per material category.

are now publicly available[2]. Each BTF material is measured in 81 illumination and 81 camera positions as an RGB image, examples of material appearance under varying illumination direction are shown in Fig. 4.

In our test set, we fixed viewpoint position to be perpendicular to material surface and included images under all 81 illumination positions. It is $15 \times 81 = 1215$ images in total, all were cropped to the same size 256×256 pixels. Training images per each material were again randomly selected and the remaining images were classified using 1-NN classifier. The number of training images went from 1 to 6 and the results were averaged over 10^3 of random selections of training images. The progress of classification accuracy is shown Fig. 8, where the CAR features outperformed the alternative features for all number of training images. The performance superiority of the CAR features is especially significant for low number of training samples, which confirms robustness of the CAR features to illumination direction variation. For additional details on robustness to illumination direction see [19,17].

3.3 Amsterdam Library of Textures

In this experiment, we tested the proposed features in the recognition of materials under combination of changing illumination spectrum and direction. The images of materials are from the recently created Amsterdam Library of Textures[3] (ALOT) [2]. The ALOT is a BTF database containing an extraordinary large collection of 250 materials, each acquired with varying viewpoint and illumination positions, and one additional illumination spectrum. Most of the materials have rough surfaces, so the movement of light source changes the appearance

[2] `http://btf.cs.uni-bonn.de`

[3] `http://staff.science.uva.nl/~mark/ALOT/`

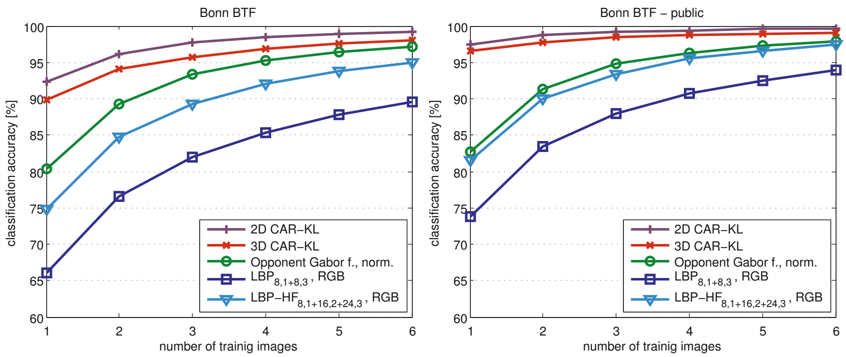

Fig. 8. Accuracy of classification on Bonn BTF database on the left and for its public part on the right. The number of training images is changing from 1 to 6 per material.

of materials. Moreover, the significant height variation of some materials (e.g. leaves) causes large and variable cast shadows, which make the recognition even more difficult.

In the "part a" of the experiment, we used one half of the dataset [2] with excluded multiple texture rotations. It contains images of the first 200 materials divided into training and test sets (1200 images each). Let c stands for camera, l for light, i for reddish illumination. The training set is defined as $c\{1,4\}l\{1,4,8\}$ and the test set contains setups $c\{2,3\}l\{3,5\}$, $c3l2$, and $c1i$. We cropped all the images to the same size 1536×660 pixels. The classification was evaluated on the test set images, where 1-NN classifier was trained on given numbers of images per material, all randomly selected from the training set.

In the "part b", we used images of all 250 materials, with all light setups, no rotations, and cameras 1 and 3, which is 14 images per material. Training images per material were randomly selected and the others were classified with 1-NN classifier, the results were averaged over 10^3 of random selections of training images. This test was performed separately for images from camera 1 and 3, the results were averaged (2×1750 images in total). As a consequence this part do not include recognition under viewpoint variation, which is in contrast with the "part a".

The results for both parts are displayed in Fig. 9, which shows the progress for different numbers of training images. The totally different scales of classification results are caused by images under different viewpoints included in the "part a" and the fact that none of the tested features are invariant to perspective projection. The viewpoint differences are even more extreme in the test set than in the training set. On the other hand, almost perfect results for 6 training images in the "part b" are not surprising, because 6 training images are leave-one-out methodology, which provides an upper bound on classification accuracy. In "part a", the CAR features outperformed the alternatives by 10% margin for all numbers of training images. In "part b", the performance of the CAR features is significantly better for low number of training images, while for leave-one-out

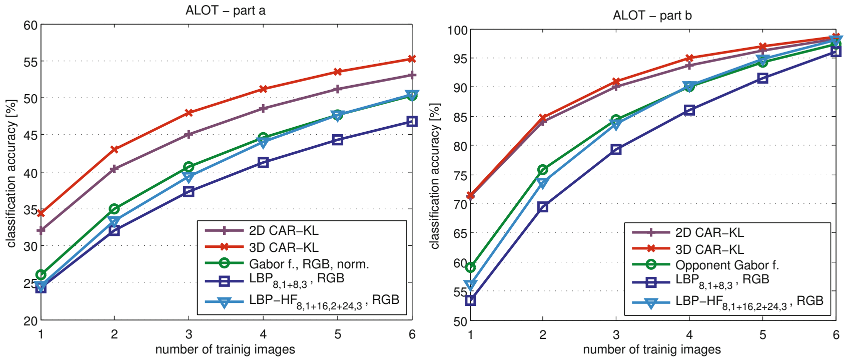

Fig. 9. Accuracy of classification on ALOT database, with the number of training images changing from 1 to 6 per material. It is worth to note different result scales caused by distinct difficulty of the setups.

the difference is about 1-2%. The *Greyworld* normalisation had minor effect on performance of Gabor and Opponent Gabor features.

The CAR features were computed on $K = 5$ pyramid levels, the results for 2D CAR with $K = 4$ went from 25.6 to 44.3 for "part a" and from 65.6 to 98.0 for "part b", which outperformed alternatives in "part b" and performed only slightly worse than Gabor features in "part a".

3.4 KTH-TIPS2 Database

Finally, the last experiment compares the performance of the proposed features on the KTH-TIPS2 database[4] [3], which includes material samples with different scales and rotations. However, as the training set always includes these scales and rotations, such an invariance is not an issue. The KTH-TIPS2 database contains 4 samples of 11 materials categories, each sample consists of images with 4 different illuminations, 3 in-plane rotations and 9 scales. The illumination conditions consist in 3 different directions plus 1 image with different spectrum. There are 4572 images and their resolution is varying around 200×200 pixels.

Training samples per each material category were again randomly selected and the remaining images were classified using 1-NN classifier. In this dataset, one training sample contains $4 \times 3 \times 9$ images, and we used from 1 to 3 samples per material category. Finally, the results were averaged over 10^3 of random selections of training samples.

Fig. 7 depicts the results, where all displayed features performed comparably. The reason is that each training sample includes images with all illumination variation, so any such invariance either do not matter or even may weaken discrimination of features. That is the reason, why Gabor features performed better without *Greyworld* normalisation.

[4] `http://www.nada.kth.se/cvap/databases/kth-tips/`

4 Conclusion

The extensive experimental evaluations illustrated in the paper prove the outstanding reliable and robust performance of efficient CAR illumination invariant Markovian textural features. The superiority of these features over leading alternatives as Gabor, Opponent Gabor, LBP, and LBP-HF features, was verified on the recent best available textural databases ALOT, Bonn BTF, UEA Uncalibrated, and KTH-TIPS2. These textural databases represent the majority of possible physically realistic acquisition conditions of surface materials represented in the form of visual textures. The proposed CAR features particularly excels in recognition with a low number of training samples, and they enable robust texture recognition in variable condition even with a single training image per material.

The results of the invariant texture retrieval or recognition can be checked online in our interactive demonstrations[5].

Acknowledgements. This research was supported by grant GAČR 102/08/0593 and partially by the projects MŠMT grants 1M0572 DAR, GAČR 103/11/0335, and CESNET 387/2010.

References

1. Ahonen, T., Matas, J., He, C., Pietikäinen, M.: Rotation Invariant Image Description with Local Binary Pattern Histogram Fourier Features. In: Salberg, A.-B., Hardeberg, J.Y., Jenssen, R. (eds.) SCIA 2009. LNCS, vol. 5575, pp. 61–70. Springer, Heidelberg (2009)
2. Burghouts, G.J., Geusebroek, J.M.: Material-specific adaptation of color invariant features. Pattern Recognition Letters 30, 306–313 (2009)
3. Caputo, B., Hayman, E., Mallikarjuna, P.: Class-specific material categorisation. In: Proceedings of the 10th IEEE International Conference on Computer Vision, ICCV 2005, October 17-21, pp. 1597–1604. IEEE (2005)
4. Chen, J., Shan, S., He, C., Zhao, G., Pietikäinen, M., Chen, X., Gao, W.: Wld: A robust local image descriptor. IEEE Transactions on Pattern Analysis and Machine Intelligence 32(9), 1705–1720 (2010)
5. Drbohlav, O., Chantler, M.: Illumination-invariant texture classification using single training images. In: Proceedings of the 4th International Workshop on Texture Analysis and Synthesis, Texture 2005, pp. 31–36 (2005)
6. Finlayson, G., Schaefer, G., Tian, G.: The UEA uncalibrated colour image database. Tech. Rep. SYS-C00-07, School of Information Systems, University of East Anglia, Norwich, United Kingdom (2000)
7. Finlayson, G., Hordley, S., Lu, C., Drew, M.: On the removal of shadows from images. IEEE Transactions on Pattern Analysis and Machine Intelligence 28(1), 59–68 (2006)
8. Finlyason, G., Xu, R.: Illuminant and gamma comprehensive normalisation in logRGB space. Patterm Recognition Letters 24, 1679–1690 (2002)

[5] `http://cbir.utia.cas.cz`

9. Haindl, M., Šimberová, S.: A Multispectral Image Line Reconstruction Method. In: Theory & Applications of Image Analysis, pp. 306–315. World Scientific Publishing Co., Singapore (1992)
10. Jain, A., Healey, G.: A multiscale representation including opponent colour features for texture recognition. IEEE Transactions on Image Processing 7(1), 124–128 (1998)
11. Mäenpää, T., Pietikäinen, M.: Classification with color and texture: jointly or separately? Pattern Recognition 37(8), 1629–1640 (2004)
12. Manjunath, B.S., Ma, W.Y.: Texture features for browsing and retrieval of image data. IEEE Transactions on Pattern Analysis and Machine Intelligence 18(8), 837–842 (1996)
13. Meseth, J., Müller, G., Klein, R.: Preserving realism in real-time rendering of bidirectional texture functions. In: OpenSG Symposium 2003, pp. 89–96. Eurographics Association, Switzerland (2003)
14. Ojala, T., Pietikäinen, M., Mäenpää, 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)
15. Santini, S., Jain, R.: Similarity measures. IEEE Transactions on Pattern Analysis and Machine Intelligence 21(9), 871–883 (1999)
16. Suen, P.H., Healey, G.: The analysis and recognition of real-world textures in three dimensions. IEEE Transactions on Pattern Analysis and Machine Intelligence 22(5), 491–503 (2000)
17. Vácha, P.: Query by pictorial example. Ph.D. thesis, Charles University in Prague, Prague (2011), http://cbir.utia.cas.cz/homepage/projects/phd_thesis/vacha_phd.pdf
18. Vacha, P., Haindl, M.: Image retrieval measures based on illumination invariant textural MRF features. In: Sebe, N., Worring, M. (eds.) Proceedings of ACM International Conference on Image and Video Retrieval, CIVR 2007, July 9-11, pp. 448–454. ACM (2007)
19. Vacha, P., Haindl, M.: Illumination invariants based on Markov random fields. In: Lovell, B., Laurendeau, D., Duin, R. (eds.) Proceedings of the 19th International Conference on Pattern Recognition, ICPR 2008, December 8-11, pp. 1–4. IEEE (2008)
20. Varma, M., Zisserman, A.: A statistical approach to texture classification from single images. International Journal of Computer Vision 62(1-2), 61–81 (2005)