Fast Features Invariant to Rotation and Scale of Texture

Milan Sulc and Jiri Matas

Center for Machine Perception,
Department of Cybernetics,
Faculty of Electrical Engineering,
Czech Technical University in Prague

Abstract. A family of novel texture representations called Ffirst, the Fast Features Invariant to Rotation and Scale of Texture, is introduced. New rotation invariants are proposed, extending the LBP-HF features, improving the recognition accuracy. Using the full set of LBP features, as opposed to uniform only, leads to further improvement. Linear Support Vector Machines with an approximate χ^2-kernel map are used for fast and precise classification.
Experimental results show that Ffirst exceeds the best reported results in texture classification on three difficult texture datasets KTH-TIPS2a, KTH-TIPS2b and ALOT, achieving 88%, 76% and 96% accuracy respectively. The recognition rates are above 99% on standard texture datasets KTH-TIPS, Brodatz32, UIUCTex, UMD, CUReT.

Keywords: Texture,classification,LBP,LBP-HF,histogram,SVM,feature maps,Ffirst

1 Introduction

Texture description and recognition techniques have been the subject to many studies for their wide range of applications. The early work focused on the problem of terrain analysis [12,36] and material inspection [37]. Later applications of texture analysis include face recognition [1], facial expressions [30,42] and object recognition [39]. The relation between scene identification and texture recognition is discussed by Renninger and Malik [28]. Texture analysis is a standard problem with several surveys available, e.g. [6,19,24,40]. Many texture description methods are based on the Local Binary Patterns [10,11,17,20–23,41], which is a computationally simple and powerful approach.

We introduce a family of novel texture representations called Ffirst - the Fast Features Invariant to Rotation and Scale of Texture. It is based on LBP-HF-S-M, the rotation invariant features obtained from sign- and magnitude-LBP histograms using Fourier transform proposed by Zhao et al. [41]. We enrich the LBP-HF-S-M representation by proposing additional rotational invariants and by the use of non-uniform patterns.

The scale invariance of Ffirst is obtained by the technique recently applied in the context of bark recognition [32].

We show that the novelties improve performance in texture recognition experiments with a feature-mapped linear SVM classifier approximating the χ^2 kernel.

The rest of this paper is organized as follows: The state-of-the-art approaches to texture recognition are briefly reviewed in Section 2. The new family of texture representations called Ffirst is introduced and described in Section 3. Section 4 is dedicated to the proposed extensions of LBP-HF and Ffirst. Section 5 presents our experiments on standard texture datasets. Section 6 concludes the paper.

2 State of the art

Several recent approaches to texture recognition report fine results on the standard datasets, often using complex description methods. Sifre and Mallat [31] used a cascade of invariants computed using scattering transforms to construct an affine invariant texture representation. A sparse representation based Earth Mover's Distance (SR-EMD) presented by Li et al. [15] achieves good results in both image retrieval and texture recognition. Quan et al. [27] propose a texture feature constructed by concatenating the lacunarity-related parameters estimated from the multi-scale local binary patterns. Local Higher-Order Statistics (LHS) proposed by Sharma et al. [30] describe higher-order differential statistics of local non-binarized pixel patterns. The method by Cimpoi et al. [7] uses Improved Fisher Vectors (IFV) for texture description. This work also shows further improvement when combined with describable texture attributes learned on the Describable Textures Dataset (DTD).

3 The Ffirst method

In order to describe texture independently of the pattern size and orientation in the image, a description invariant to rotation and scale is needed. For practical applications we also demand computational efficiency.

In this section we introduce a new texture description called Ffirst (Fast Features Invariant to Rotation and Scale of Texture), which combines several state-of-the-art approaches to satisfy the given requirements. This method builds on and improves a texture descriptor for bark recognition introduced in [32].

3.1 Completed Local Binary Pattern and Histogram Fourier Features

The Ffirst description is based on the Local Binary Patterns (LBP) [20,22]. The common LBP operator (further denoted as sign-LBP) computes the signs of differences between pixels in the 3×3 neighbourhood and the center pixel. LBP have been generalized [21] to arbitrary number of neighbours P on a circle of

radius R, using an image function $f(x, y)$ and neighbourhood point coordinates (x_p, y_p):

$$\mathrm{LBP}_{P,R}(x,y) = \sum_{p=0}^{P-1} s(f(x,y) - f(x_p,y_p))2^p, \quad s(z) = \begin{cases} 1 : z \leq 0 \\ 0 : \quad \text{else} \end{cases}. \tag{1}$$

To achieve rotation invariance[1], Ffirst uses the so called LBP Histogram Fourier Features (LBP-HF) introduced by Ahonen et al. [2], which describe the histogram of uniform patterns using coefficients of the discrete Fourier transform. Uniform LBP are patterns with at most 2 spatial transitions (bitwise 0-1 changes). Unlike the simple rotation invariants using $\mathrm{LBP}^{\mathrm{ri}}$ [21,25], which assign all uniform patterns with the same number of 1s into one bin,

$$\mathrm{LBP}^{ri}_{P,R} = \min\left\{\mathrm{ROR}\left(\mathrm{LBP}_{P,R}, i\right) \mid i = 0, 1, .., P-1\right\}, \tag{2}$$

the LBP-HF features preserve the information about relative rotation of the patterns.

Denoting a uniform pattern $U_p^{n,r}$, where n is the "orbit" number corresponding to the number of "1" bits and r denotes the rotation of the pattern, the DFT for given n is expressed as:

$$H(n,u) = \sum_{r=0}^{P-1} h_I(U_p^{n,r}) e^{-i2\pi ur/P}, \tag{3}$$

where the histogram value $h_I(U_p^{n,r})$ denotes the number of occurrences of a given uniform pattern in the image.

The LBP-HF features are equal to the absolute value of the DFT magnitudes (which are not influenced by the phase shift caused by rotation):

$$\text{LBP-HF}(n,u) = |H(n,u)| = \sqrt{H(n,u)\overline{H(n,u)}}. \tag{4}$$

Since h_I are real, $H(n,u) = H(n,P-u)$ for $u = (1, .., P-1)$, and therefore only $\left\lfloor \frac{P}{2} \right\rfloor + 1$ of the DFT magnitudes are used for each set of uniform patterns with n "1" bits for $0 < n < P$. Three other bins are added to the resulting representation, namely two for the "1-uniform" patterns (with all bins of the same value) and one for all non-uniform patterns.

The LBP histogram Fourier features can be generalized to any set of uniform patterns. In Ffirst, the LBP-HF-S-M description introduced by Zhao et al. [41] is used, where the histogram Fourier features of both sign- and magnitude-LBP are calculated to build the descriptor. The combination of both sign- and magnitude-LBP called Completed Local Binary Patterns (CLBP) was introduced by Guo and Zhang [10]. The magnitude-LBP checks if the magnitude of the difference of

[1] LBP-HF (as well as LBP^{ri}) are rotation invariant only in the sense of a circular bit-wise shift, e.g. rotation by multiples 22.5° for $\mathrm{LBP}_{16,R}$.

the neighbouring pixel (x_p, y_p) against the central pixel (x, y) exceeds a threshold t_p:

$$\text{LBP-M}_{P,R}(x,y) = \sum_{p=0}^{P-1} s(|f(x,y) - f(x_p,y_p)| - t_p)2^p. \tag{5}$$

We adopted the common practice of choosing the threshold value (for neighbours at p-th bit) as the mean value of all m absolute differences in the whole image:

$$t_p = \sum_{i=1}^{m} \frac{|f(x_i,y_i) - f(x_{ip},y_{ip})|}{m}. \tag{6}$$

The LBP-HF-S-M histogram is created by concatenating histograms of LBP-HF-S and LBP-HF-M (computed from uniform sign-LBP and magnitude-LBP).

3.2 Multi-scale description and scale invariance

A scale space is built by computing LBP-HF-S-M from circular neighbourhoods with exponentially growing radius R. Gaussian filtering is used[2] to overcome noise.

Unlike the MS-LBP approach of Mäenpää and Pietikäinen [17], where the radii of the LBP operators are chosen so that the effective areas of different scales touch each other, Ffirst uses a finer scaling with a $\sqrt{2}$ step between scales radii R_i, i.e. $R_i = R_{i-1}\sqrt{2}$.

This radius change is equivalent to decreasing the image area to one half. The finer sampling uses more evenly spaced information compared to [17], as illustrated in Figures 1a, 1b. The first LBP radius used is $R_1 = 1$, as the LBP with low radii capture important high frequency texture characteristics.

Similarly to [17], the filters are designed so that most of their mass lies within an effective area of radius r_i. We select the effective area diameter, such that the effective areas at the same scale touch each other: $r_i = R_i \sin \frac{\pi}{P}$.

LBP-HF-S-M histograms from c adjacent scales are concatenated into a single descriptor. Invariance to scale changes is increased by creating n_{conc} multi-scale descriptors for one image. See Algorithm 1 for the overview of the texture description method.

[2] The Gaussian filtering is used for a scale i only if $\sigma_i > 0.6$, as filtering with lower σ_i leads to significant loss of information.

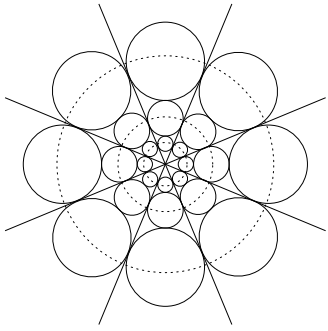

(a) Scale space of Mäenpää and Pietikäinen [17]

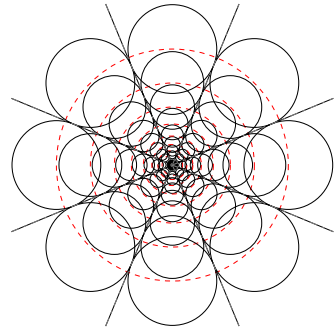

(b) Scale space from [32] used in Ffirst

Fig. 1. The effective areas of filtered pixel samples in a multi-resolution $\mathrm{LBP}_{8,R}$ operator

Algorithm 1 The Ffirst description method overview

```
R_1 := 1
for all scales i = 1...(n_conc + c − 1) do
  σ_i := R_i sin(π/P)/1.96
  if σ_i > 0.6 then
    apply Gaussian filter (with std. dev. σ_i) on the original image
  end if
  extract LBP_{P,R_i}-S and LBP_{P,R_i}-M and build the LBP-HF-S-M descriptor
  for j = 1...c do
    if i ≥ j and i < j + n_conc then
      attach the LBP-HF-S-M to the j-th multi-scale descriptor
    end if
  end for
  R_{i+1} := R_i √2
end for
```

3.3 Support Vector Machine and feature maps

In most applications, a Support Vector Machine (SVM) classifier with a suitable non-linear kernel provides higher recognition accuracy at the price of significantly higher time complexity and higher storage demands (dependent on the number of support vectors). An approach for efficient use of additive kernels via explicit feature maps is described by Vedaldi and Zisserman [35] and can be combined with a linear SVM classifier. Using linear SVMs on feature-mapped data improves the recognition accuracy, while preserving linear SVM advantages like fast evaluation and low storage (independent on the number of support vectors), which are both very practical in real time applications. In Ffirst we use the explicit feature map approximation of the χ^2 kernel.

The "One versus All" classification scheme is used for multi-class classification, implementing the Platt's probabilistic output [16,26] to ensure SVM results comparability among classes. The maximal posterior probability estimate over all scales is used to determine the resulting class.

In our experiments we use a Stochastic Dual Coordinate Ascent [29] linear SVM solver implemented in the VLFeat library [34].

4 Adding rotational invariants

The LBP-HF features used in the proposed Ffirst description are built from the DFT magnitudes of differently rotated uniform patterns, as described in Section 3.1. We propose 3 more variants for the description, which will appear in our experiments in Section 5.

The variant denoted as Ffirst$^+$ creates additional rotational invariants, LBP-HF$^+$ features, computed from the first harmonics for each orbit:

$$\text{LBP-HF}^+(n) = \sqrt{H(n,1)\overline{H(n+1,1)}} \tag{7}$$

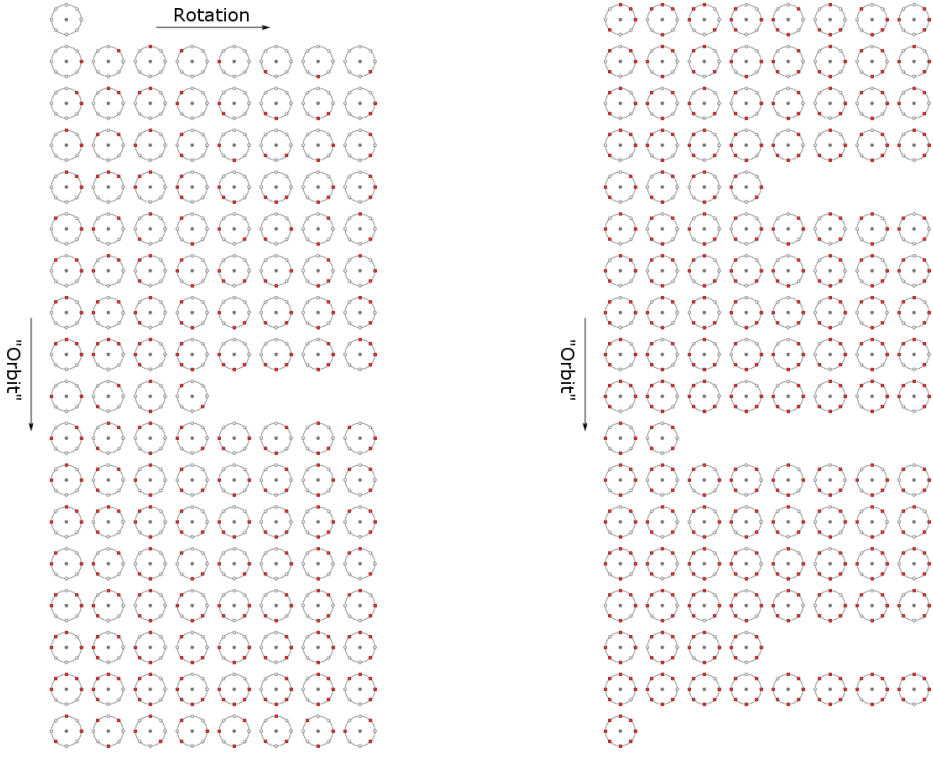

Fig. 2. Ordering the full set of Local Binary Patterns for the Histogram Fourier features

Another variant, Ffirst$^{\forall}$, uses all LBP instead of only the subset of uniform patterns. Note that in this case, some orbits have a lower number of patterns, as some non-uniform patterns have less possible rotations, as illustrated in Figure 2.

The last variant, denoted as Ffirst$^{\forall+}$, uses the full set of patterns for LBP-HF features, adding also the additional LBP-HF^{+} features

5 Experiments

5.1 Datasets

The proposed Ffirst method for texture classification was tested using the standard evaluation protocols on the following texture datasets:

The KTH-TIPS texture database [9, 13] contains images of 10 materials. There are 81 images (200x200 px) of each material with different combination of pose, illumination and scale.

The standard evaluation protocol on the KTH-TIPS dataset uses 40 training images per material.

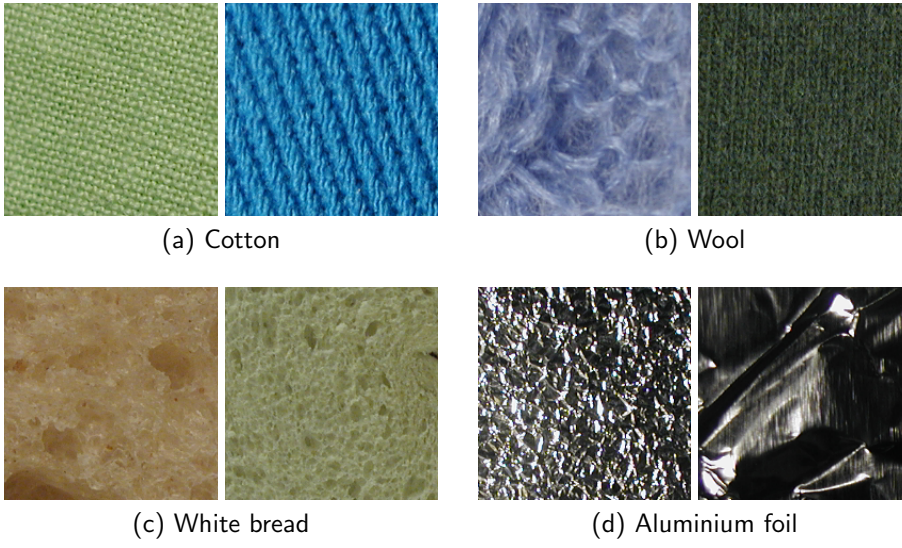

Fig. 3. Examples of 4 texture classes from the KTH-TIPS2 database

The KTH-TIPS2 database was published [5, 18] shortly after KTH-TIPS. It builds on the KTH-TIPS database, but provides multiple sets of images - denoted as "samples" - per material class (examples in Figure 3).

There are 4 "samples" for each of the 11 materials in the KTH-TIPS2 database, containing 108 images per "sample" (again with different combination of pose, illumination and scale). However, in the first version of this dataset, for 4 of those 44 "samples" only 72 images were used. This first version is usually denoted as KTH-TIPSa, and the standard evaluation method uses 3 "samples" from each class for training and 1 for testing. The "complete" version of this database, KTH-TIPSb, is usually trained only on 1 "samples" per class and tested on the remaining 3 "samples".

The Brodatz32 dataset [33] was published in 1998 and it contains low resolution (64x64 px) grey-scale images of 32 textures from the photographs published by Phil Brodatz [3] in 1966, with artificially added rotation (90°) and scale change (a 64x64 px scaled block obtained from 45x45 pixels in the middle). There are 64 images for each texture class in total.

The standard protocol for this dataset simply divides the data into two halves (i.e. 32 images per class in the training set and 32 in the test set).

Even though the original images are copyrighted and the legality of their usage in academic publications is unclear[3], Brodatz textures are one of the most popular and broadly used sets in texture analysis.

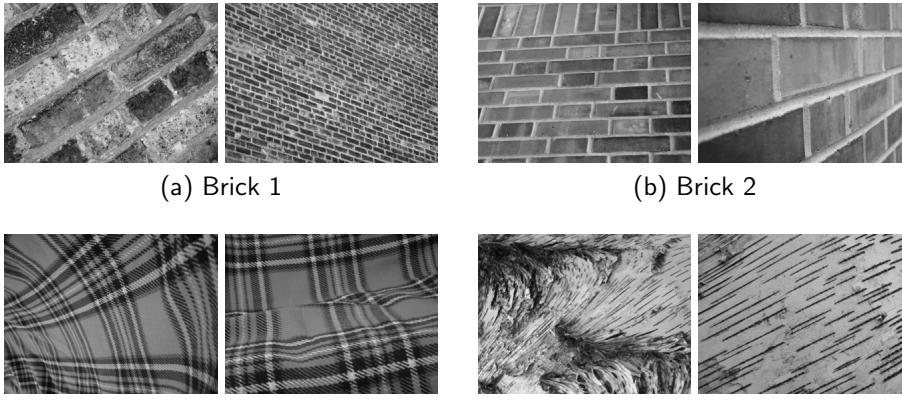

Fig. 4. Examples of 4 texture classes from the UIUCTex database

The UIUCTex database, sometimes referred to as the Ponce Group Texture Database, was published by Lazebnik et al. [14] in 2005 and features 25 different texture classes, 40 samples each. All images are in VGA resolution (640x480 px) and in grey-scale.

[3] `http://graphics.stanford.edu/projects/texture/faq/brodatz.html`

The surfaces included in the database are of various nature (wood, marble, gravel, fur, carpet, brick, ..) and were acquired with significant viewpoint, scale and illumination changes and additional sources of variability, including, but not limited to, non-rigid material deformations (fur, fabric, and water) and viewpoint-dependent appearance variations (glass). Examples of images from different classes are in Figure 4.

The results on this dataset are usually evaluated using 20 or 10 training images per class. In our experiments, the former case with a larger training set is performed.

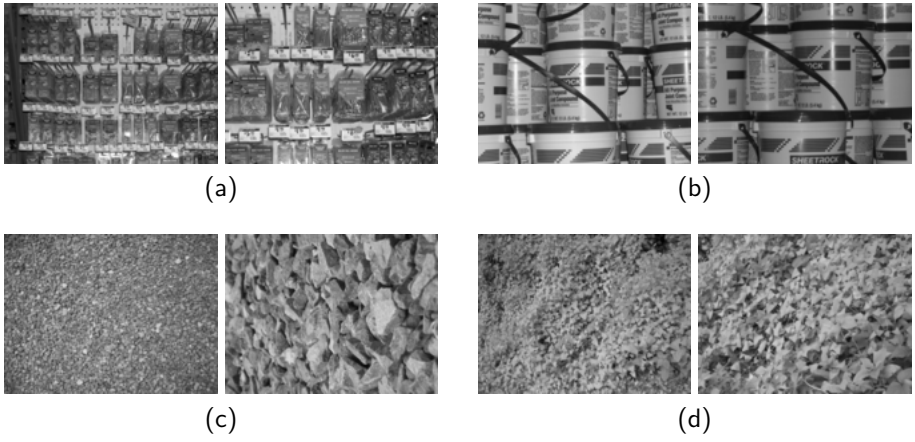

Fig. 5. Examples of 4 texture classes from the UMD database

The UMD dataset [38] consists of 1000 uncalibrated, unregistered grey-scale images of size 1280x960 px, 40 images for each of 25 different textures. The UMD database contains non-traditional textures like images of fruits, shelves of bottles and buckets, various plants, or floor textures.

The standard evaluation protocol for UMD is dividing the data into two halves (i.e. 20 images per class in the training set and 20 in the test set).

The CUReT image database [8] contains textures from 61 classes, each observed with 205 different combinations of viewing and illumination directions. In the commonly used version, denoted as the cropped CUReT database[4], only 92 images are chosen, for which a sufficiently large region of texture is visible across all materials. A central 200x200 px region is cropped from each of these images, discarding the remaining background. There are thus 61x92=5612 images in the cropped database.

[4] http://www.robots.ox.ac.uk/ vgg/research/texclass/setup.html

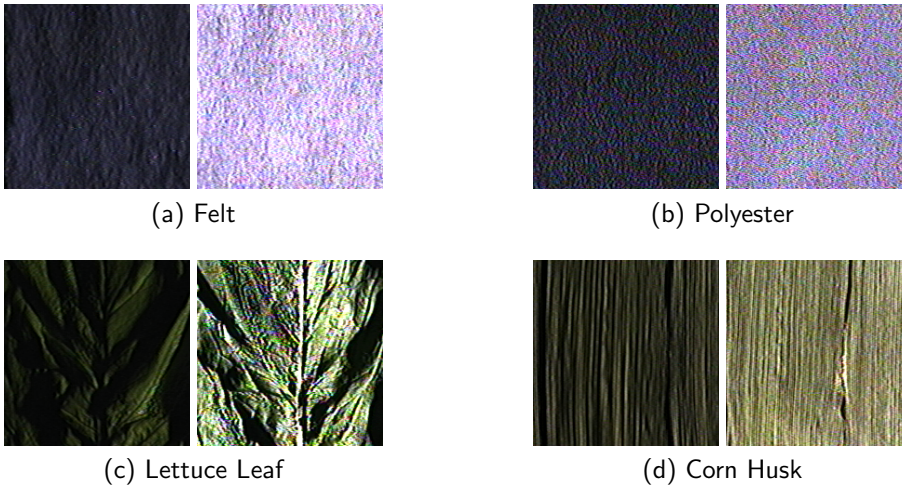

Fig. 6. Examples of 4 texture classes from the CUReT database

Though CUReT also contains a BRDF (bidirectional reflectance distribution function) database, for purposes of standard texture recognition methods, only the image database is used. We use 46 training images per class, which is a standard evaluation protocol for the CUReT database.

Fig. 7. Examples of 4 texture classes from the ALOT database

The Amsterdam Library of Textures [4], denoted as ALOT, contains 250 texture classes. Each class contains 100 images obtained with different combinations of viewing and illumination directions and illumination color.

To compare our results on the ALOT dataset to the state-of-the-art [27] we use 20 training images and 80 test images per class.

5.2 Parameter setting

In all following experiments, we use the same setting of our method: $n_{\text{conc}} = 3$ multi-scale descriptors per image are used, each of them consisting of $c = 6$ scales described using LBP-HF-S-M. A higher number of concatenated scales offers only minimal improvement in accuracy, while increasing the processing time. The final histogram is kernelized using the approximate χ^2 feature map, although using the intersection kernel would provide similar results. In the application, the data are only trained once and the training precision is more important than the training time. Thus we demand high accuracy, setting SVM parameters to: regularization parameter $\lambda = 10^{-7}$, tolerance for the stopping criterion $\epsilon = 10^{-7}$, maximum number of iterations: 10^8. We use the unified setting in order to show the generality of the Ffirst description, although setting the parameters individually for a given dataset might further increase the accuracy.

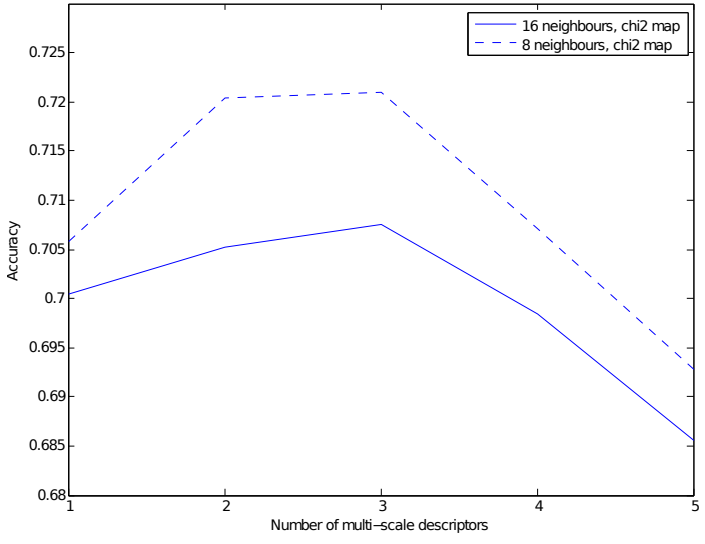

Fig. 8. Dependence of the KTH-TIPS2b recognition rate on the number of multiscale descriptors in Ffirst, denoted c.

Figures 8 and 9 illustrate the effect of different parameter settings on the recognition accuracy for the KTH-TIPS2b texture database.

To reduce the effect of random training and test data choice, the presented results are averaged from 10 experiments.

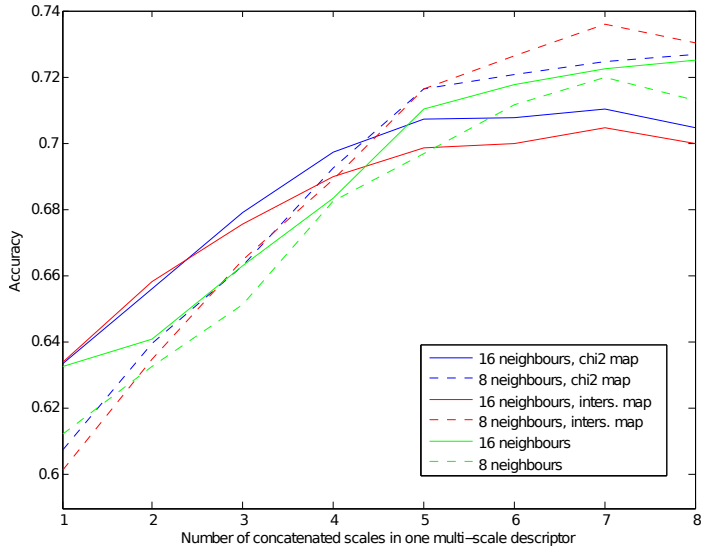

Fig. 9. Feature mapping and concatenating features from multiple scales in Ffirst, KTH-TIPS2b

5.3 Classification results

Table 1. Evaluation of Ffirst on other standard datasets, compared to the state-of-the-art methods

	Brodatz32	*UIUCTex*	*UMD*	*CUReT*	*ALOT*
Num. of classes	32	25	25	61	250
Ffirst	99.2±0.3	98.6±0.6	99.3±0.3	98.5±0.2	92.9±0.3
Ffirst^{+}	99.3±0.3	98.7±0.7	99.3±0.3	98.6±0.3	93.4±0.3
Ffirst$^{\forall}$	99.6±0.2	99.0±0.5	99.3±0.3	99.1±0.2	95.0±0.3
Ffirst$^{\forall+}$	**99.7±0.2**	99.3±0.4	99.3±0.3	99.2±0.2	**95.9±0.5**
IFV_{SIFT} [7]	–	97.0±0.9	99.2±0.4	99.6±0.3	–
IFV_{SIFT} [7] + DeCAF [5]	–	99.0±0.5	99.5±0.3	**99.8±0.2**	–
Scattering [31]	–	**99.4±0.4**	**99.7±0.3**	–	–
LHS [30]	99.5±0.2	–	–	–	–
SR-EMD-M [15]	–	–	**99.9**	99.5	–
PLS [27]	–	96.6	98.99	–	93.4
MS-LBP-HF-KlSVM [32]	96.2±0.6	96.4±0.6	–	–	–

[5] Results from `http://www.robots.ox.ac.uk/~vgg/data/dtd/`

The experimental results in texture classification are compared to the state-of-the-art in Tables 1,2, containing the results on the KTH-TIPS datasets and on other standart texture datasets respectively.

Table 2. Evaluation of Ffirst on the KTH-TIPS datasets, compared to the state-of-the-art methods

	KTH-TIPS2a	*KTH-TIPS2b*	*KTH-TIPS*
Num. of classes	11	11	10
Ffirst	86.2±5.5	72.1±5.1	98.9±0.7
Ffirst^{+}	86.4±5.0	72.7±5.2	98.9±0.8
Ffirst$^{\forall}$	88.0±6.5	75.8±4.1	99.1±0.5
Ffirst$^{\forall+}$	**88.2±6.7**	**76.0±4.1**	99.1±0.5
IFV_{SIFT} [7]	82.5±5.2	69.3±1.0	**99.7±0.1**
IFV_{SIFT} [7] + DeCAF [6]	84.4±1.8	**76.0±2.9**	99.8±0.2
IFV_{SIFT} [7] + DeCAF +DTD_{RBF} [6][7]	–	77.4±2.2	–
Scattering [31]	–	–	99.4±0.4
LHS [30]	73.0±4.7	–	–
SR-EMD-M [15]	–	–	**99.8**
PLS [27]	–	–	98.4

[6] Results from `http://www.robots.ox.ac.uk/~vgg/data/dtd/`
[7] The method requires an additional training set (the DTD dataset)

5.4 Suitability for real-time applications

Table 3 shows a comparison of our image processing times to the state-of-the-art texture recognition method by Cimpoi et al. [7] based on IVF_{SIFT}. Both the implementation of Ffirst and IVF_{SIFT}[8] used MATLAB scripts with a C code in the VLFeat [34] framework (after adding a new CLBP implementation for our method). The processing times were measured on a standard laptop (1.3 GHz Intel Core i5, 4 GB 1600 MHz DDR3) without parallelization.

The average description time for a low resolution (200x200px) image for Ffirst is at most 0.05 s, while for higher resolutions the processing time will grow proportionally to the image resolution, as the number of local operations will increase with the number of pixels.

[8] Using the code kindly provided by the authors of [7]

Table 3. Average image description time for one image, compared to IFV_{SIFT}

	KTH-TIPS2b	*KTH-TIPS*	*CUReT*
Image resolution	200x200 px	200x200 px	200x200 px
Ffirst	0.029 s / im.	0.028 s / im.	0.029 s / im.
$Ffirst^{+}$	0.032 s / im.	0.032 s / im.	0.032 s / im.
$Ffirst^{\forall}$	0.035 s / im.	0.035 s / im.	0.036 s / im.
$Ffirst^{\forall+}$	0.048 s / im.	0.049 s / im.	0.049 s / im.
IFV_{SIFT} [7]	0.089 s / im.	0.088 s / im.	0.090 s / im.

6 Conclusions

We proposed a family of novel texture representations called Ffirst, the Fast Features Invariant to Rotation and Scale of Texture, using several state-of-the-art approaches. The first variant, $Ffirst^{+}$, uses newly proposed rotational invariants, another, denoted as $Ffirst^{\forall}$, allows to build the features from the full set of LBP, including non-uniform patterns.

The $Ffirst^{\forall+}$ method, using both proposed improvements, achieves the best results, exceeding the best reported results in texture classification on three difficult texture datasets, KTH-TIPS2a, KTH-TIPS2b and ALOT, achieving 88%, 76% and 96% accuracy respectively. The recognition rates were above 99% on standard texture datasets KTH-TIPS, Brodatz32, UIUCTex, UMD, CUReT.

The Ffirst description and the evaluation based on linear Support Vector Machines are fast, making the proposed method suitable for real time applications.

Acknowledgement

Jiri Matas was supported by Czech Science Foundation Project GACR P103/12/G084, Milan Sulc by Czech Technical University project SGS13/142/OHK3/2T/13.

References

1. Ahonen, T., Hadid, A., Pietikainen, M.: Face description with local binary patterns: Application to face recognition. Pattern Analysis and Machine Intelligence, IEEE Transactions on 28(12), 2037–2041 (2006)
2. Ahonen, T., Matas, J., He, C., Pietikäinen, M.: Rotation invariant image description with local binary pattern histogram fourier features. In: SCIA '09, in Proc. pp. 61–70. Springer-Verlag (2009)
3. Brodatz, P.: Textures: a photographic album for artists and designers, vol. 66. Dover New York (1966)
4. Burghouts, G.J., Geusebroek, J.M.: Material-specific adaptation of color invariant features. Pattern Recognition Letters 30(3), 306–313 (2009)

5. Caputo, B., Hayman, E., Mallikarjuna, P.: Class-specific material categorisation. In: Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on. vol. 2, pp. 1597–1604. IEEE (2005)
6. Chen, C.h., Pau, L.F., Wang, P.S.p.: Handbook of pattern recognition and computer vision. World Scientific (2010)
7. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. arXiv preprint arXiv:1311.3618 (2013)
8. Dana, K.J., Van Ginneken, B., Nayar, S.K., Koenderink, J.J.: Reflectance and texture of real-world surfaces. ACM Transactions on Graphics (TOG) 18(1), 1–34 (1999)
9. Fritz, M., Hayman, E., Caputo, B., Eklundh, J.O.: The kth-tips database (2004)
10. Guo, Z., Zhang, D.: A completed modeling of local binary pattern operator for texture classification. Image Processing, IEEE Transactions on 19(6), 1657–1663 (2010)
11. Guo, Z., Zhang, L., Zhang, D.: Rotation invariant texture classification using lbp variance (lbpv) with global matching. Pattern recognition 43(3), 706–719 (2010)
12. Haralick, R.M., Shanmugam, K., Dinstein, I.H.: Textural features for image classification. Systems, Man and Cybernetics, IEEE Transactions on (6), 610–621 (1973)
13. Hayman, E., Caputo, B., Fritz, M., Eklundh, J.O.: On the significance of real-world conditions for material classification. In: Computer Vision-ECCV 2004, pp. 253–266. Springer (2004)
14. Lazebnik, S., Schmid, C., Ponce, J.: A sparse texture representation using local affine regions. PAMI 27(8), 1265–1278 (2005)
15. Li, P., Wang, Q., Zhang, L.: A novel earth movers distance methodology for image matching with gaussian mixture models. In: Proc. of IEEE International Conference on Computer Vision (ICCV) (2013)
16. Lin, H.T., Lin, C.J., Weng, R.C.: A note on platts probabilistic outputs for support vector machines. Machine learning 68(3) (2007)
17. Mäenpää, T., Pietikäinen, M.: Multi-scale binary patterns for texture analysis. In: Image Analysis, pp. 885–892. Springer (2003)
18. Mallikarjuna, P., Fritz, M., Targhi, A., Hayman, E., Caputo, B., Eklundh, J.: The kth-tips and kth-tips2 databases. `http://www.nada.kth.se/cvap/databases/kth-tips` (2006)
19. Mirmehdi, M., Xie, X., Suri, J.: Handbook of texture analysis. Imperial College Press (2009)
20. Ojala, T., Pietikainen, M., Harwood, D.: Performance evaluation of texture measures with classification based on kullback discrimination of distributions. In: IAPR 1994, in Proc. vol. 1, pp. 582–585 vol.1 (1994)
21. Ojala, T., Pietikainen, M., Maenpaa, T.: Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. PAMI 24(7), 971–987 (2002)
22. Ojala, T., Pietikäinen, M., Harwood, D.: A comparative study of texture measures with classification based on featured distributions. Pattern Recognition 29(1), 51–59 (1996)
23. Ojala, T., Valkealahti, K., Oja, E., Pietikäinen, M.: Texture discrimination with multidimensional distributions of signed gray-level differences. Pattern Recognition 34(3), 727–739 (2001)
24. Pietikäinen, M.: Texture recognition. Computer Vision: A Reference Guide pp. 789–793 (2014)
25. Pietikäinen, M., Ojala, T., Xu, Z.: Rotation-invariant texture classification using feature distributions. Pattern Recognition 33(1), 43–52 (2000)

26. Platt, J.: Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. Advances in large margin classifiers 10(3) (1999)
27. Quan, Y., Xu, Y., Sun, Y., Luo, Y.: Lacunarity analysis on image patterns for texture classification. In: Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on. IEEE (2014)
28. Renninger, L.W., Malik, J.: When is scene identification just texture recognition? Vision research 44(19), 2301–2311 (2004)
29. Shalev-Shwartz, S., Zhang, T.: Stochastic dual coordinate ascent methods for regularized loss minimization. arXiv preprint arXiv:1209.1873 (2012)
30. Sharma, G., ul Hussain, S., Jurie, F.: Local higher-order statistics (lhs) for texture categorization and facial analysis. In: Computer Vision–ECCV 2012, pp. 1–12. Springer (2012)
31. Sifre, L., Mallat, S.: Rotation, scaling and deformation invariant scattering for texture discrimination. In: Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on. pp. 1233–1240. IEEE (2013)
32. Sulc, M., Matas, J.: Kernel-mapped histograms of multi-scale lbps for tree bark recognition. In: Image and Vision Computing New Zealand (IVCNZ), 2013 28th International Conference of. pp. 82–87 (Nov 2013)
33. Valkealahti, K., Oja, E.: Reduced multidimensional co-occurrence histograms in texture classification. PAMI 20(1), 90–94 (1998)
34. Vedaldi, A., Fulkerson, B.: VLFeat: An open and portable library of computer vision algorithms. `http://www.vlfeat.org/` (2008)
35. Vedaldi, A., Zisserman, A.: Efficient additive kernels via explicit feature maps. PAMI 34(3) (2011)
36. Weszka, J.S., Dyer, C.R., Rosenfeld, A.: A comparative study of texture measures for terrain classification. Systems, Man and Cybernetics, IEEE Transactions on (4), 269–285 (1976)
37. Weszka, J.S., Rosenfeld, A.: An application of texture analysis to materials inspection. Pattern Recognition 8(4), 195–200 (1976)
38. Xu, Y., Ji, H., Fermüller, C.: Viewpoint invariant texture description using fractal analysis. International Journal of Computer Vision 83(1), 85–100 (2009)
39. Zhang, J., Marszałek, M., Lazebnik, S., Schmid, C.: Local features and kernels for classification of texture and object categories: A comprehensive study. IJCV 73(2), 213–238 (2007)
40. Zhang, J., Tan, T.: Brief review of invariant texture analysis methods. Pattern recognition 35(3), 735–747 (2002)
41. Zhao, G., Ahonen, T., Matas, J., Pietikainen, M.: Rotation-invariant image and video description with local binary pattern features. Image Processing, IEEE Transactions on 21(4), 1465–1477 (2012)
42. Zhao, G., Pietikainen, M.: Dynamic texture recognition using local binary patterns with an application to facial expressions. Pattern Analysis and Machine Intelligence, IEEE Transactions on 29(6), 915–928 (2007)