Pattern Recognition Letters
journal homepage: www.elsevier.com

Rotation invariant co-occurrence features based on digital circles and discrete Fourier transform

Francesco Bianconi[a], Antonio Fernández[b,**]

[a]*Department of Engineering, Università degli Studi di Perugia, Via G. Duranti 67, 06125 Perugia, Italy*
[b]*School of Industrial Engineering, Universidade de Vigo, Campus Universitario Lagoas-Marcosende, 36310 Vigo, Spain*

ARTICLE INFO

Article history:

Texture classification
co-occurrence matrices
rotation invariance
digital circles
discrete Fourier transform.

ABSTRACT

Grey-level co-occurrence matrices (GLCM) have been on the scene for almost forty years and continue to be widely used today. In this paper we present a method to improve accuracy and robustness against rotation of GLCM features for image classification. In our approach co-occurrences are computed through digital circles as an alternative to the standard four directions. We use discrete Fourier transform normalization to convert rotation dependent features into rotation invariant ones. We tested our method on four different datasets of natural and synthetic images. Experimental results show that our approach is more accurate and robust against rotation than the standard GLCM features.

1. Introduction

Grey-level co-occurrence matrices are among the most long-standing texture descriptors in use, their origin dating back to the pioneering work of Haralick et al. (1973). Though many other methods have been proposed since their appearance – see Xie and Mirmehdi (2008) for a comprehensive overview – GLCM continue to be very common and widely adopted still today. Bibliometric data reveal that the number of relevant scientific papers has even increased during the last years (see Tab. 1). GLCM features are particularly appealing for their conceptual simplicity, ease of implementation and the low number of features they produce. A recent comparative experiment on image classification under non-ideal conditions (Kandaswamy et al., 2011) showed that GLCM features tend to perform better when few classes (10 or less) are involved, a situation in which they can compete with newer and more powerful methods. Besides, co-occurrence features can be combined with other descriptors that convey complementary information through suitable fusion schemes (Clausi and Deng, 2005). Recent applications of GLCM span very diverse areas of image processing, including surface inspection (Dutta et al., 2012; Ben Salem and Nasri, 2010), environmental monitoring (Arebey et al., 2012; Manivannan et al., 2012), content-based image retrieval (Ramamurthy and Chandran, 2012) and image reconstruction (Azzabou et al., 2010). Among the numerous application areas, co-occurrence matrices seem to be particularly common in medical image analysis (Hu et al., 2012; Linder et al., 2012; Barwad et al., 2012; Parekh, 2012; Gómez et al., 2012) and remote sensing (Berthelot et al., 2013; Wu et al., 2012; Masetti and Calder, 2012; Kandaswamy et al., 2005).

Co-occurrence matrices have been extended in various directions, leading to several variations such as generalized co-occurrence matrices (Davis et al., 1979), which consider the distribution of local maxima; integrative co-occurrence matrices (Palm, 2004), which operate on colour images and, more recently, pattern co-occurrence matrices (Song, 2011; González et al., 2014), which analyse the co-occurrence of local patterns. By contrast, the original formulation has not changed significantly since its appearance. This is not uncommon: when a method matures and new ones appear, scientific interest tends to switch from the former to the latter. Newer methods receive more attention, and the older becomes frozen, somewhat immutable, with few chances of improvement. Something of this

**Corresponding author: Tel.: +34-986-818602; fax: +34-986-812201
e-mail: bianco@ieee.org (Francesco Bianconi), antfdez@uvigo.es (Antonio Fernández)
URL: http://dismac.dii.unipg.it/bianco (Francesco Bianconi), http://webs.uvigo.es/antfdez (Antonio Fernández)

type we believe has happened with co-occurrence matrices, at least for what it concerns rotation invariant features.

Motivated by the wide diffusion of the method – even in very critical areas like medical image analysis and computer-assisted diagnosis, we wished to investigate whether it was possible to improve robustness and accuracy of the method in rotation invariant classification tasks. This is a major concern, for in many applications images can occur in different and uncontrolled rotation angles. The common approach to obtaining rotationally-invariant features from co-occurrence matrices consists of averaging (Haralick et al., 1973; Alam and Faruqui, 2011; Bino Sebastian et al., 2012) or – equivalently – summing up (Petrou and García Sevilla, 2006, p. 215) the matrices corresponding to the same distance and different directions. We believe that this procedure reduces significantly and somewhat unnecessarily the discrimination capability of the resulting features. We therefore propose some improvements to compute more efficient rotationally-invariant features from GLCM. Our study considers the effects of two design factors that determine how GLCM features are computed. These are: 1) the spatial arrangement of pairs of pixels; and 2) the way to convert GLCM features into rotation invariant ones. In the remainder of the paper we first discuss such design factors (Sec. 2), then evaluate their effects through an image classification experiment (Sec. 3). We present and analyse the results in Sec. 4 and conclude with final considerations in Sec. 5.

2. Design factors

Grey-level co-occurrence matrices estimate the joint occurrence probability of grey levels at a given distance and direction. The method is intrinsically directional, hence sensitive to rotation. In order to achieve rotation invariant descriptors, we need to remove the dependence on direction and obtain features that depend on distance only. Such a goal can be obtained through the following steps: 1) for each pixel in the image, consider all pixels that are located approximately at a given distance from it (we refer to this entity as neighbourhood); 2) extract rotation dependent features for each direction defined by the neighbourhood; 3) convert the rotation dependent features into rotation-independent ones. We discuss each step in the remainder of this section, with particular emphasis on steps 1) and 3), which represent the chief objective of this study. Less emphasis will be devoted to step 2), for the extraction of rotation dependent co-occurrence features is quite a standard operation that does not require further explanations. Following the terminology of Design of Experiments (see Cochran and Cox 1957, p. 148), the type of neighbourhood and the procedure to obtain rotation invariant features will be the *design factors* of our study; the possible solutions for each factor will be referred to as *variations*[1]. A combination of variations will be referred to as a *treatment*. In the following subsections we discuss the design factors and variations considered in the study.

[1]To our ear this term sounds better, in this context, than the more common *level*.

2.1. Type of neighbourhood

Let I_0 denote the grey value of a pixel and I_j, $j \in \{1, \ldots, N\}$ the grey values of a set of pixels approximately equidistant from it, with the convention that I_1 stands horizontally on the right of I_0 and the others follow counter-clockwise from I_1 (Fig. 1). Note that, since the image is scanned by one-pixel steps, it is not necessary, due to symmetry, to consider the entire neighbourhood: only one half suffices (for a detailed explanation see Petrou and García Sevilla 2006, p. 280).

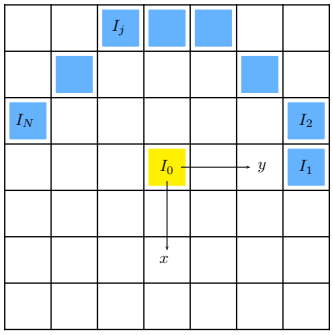

Fig. 1. Neighbourhood of pixels.

We consider two variations for the type of neighbourhood: the original formulation proposed by Haralick et al. (1973), based on four directions – this is by far the most used in practice –, and the digital circles proposed by Petrou and García Sevilla (2006).

2.1.1. Original formulation

In Haralick's formulation the neighbourhood is formed by the central pixel plus four peripheral pixels equally spaced at angular intervals of 45°. For a given distance d, the relative coordinates of the peripheral pixels with respect to the central one are: $(0, d)$, $(-d, d)$, $(-d, 0)$, $(-d, -d)$. In this scheme the number of pixels in the neighbourhood is constant, therefore independent of d. Fig. 2 shows the neighbourhoods corresponding to distances $d = 1, 2, 3$. In the remainder, we use the subscript '$*$' to refer to this variation.

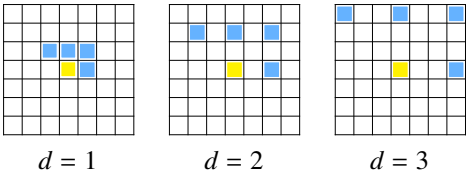

Fig. 2. Type of neighbourhood: original, four-direction.

2.1.2. Digital circles

As an alternative to the neighbourhood described above, Petrou and García Sevilla (2006) suggested the use of digital circles. The definition of circle in the continuous space does not translate immediately into the digital domain, thus the way to define digital circles is not unique (for a discussion on this

Table 1. Co-occurrence matrices – bibliometric overview. Source: Scopus®. Query: TITLE("co-occurrence matrices" OR "GLCM") OR AUTHKEY("co-occurrence matrices" OR "GLCM"). Accessed on Jan 27, 2014

Source type	Year							
	2006	2007	2008	2009	2010	2011	2012	2013
Conference proc.	34	33	48	74	95	113	103	63
Journals	28	30	35	47	45	60	67	93

topic see the work of Mukherjee et al. 2000). Herein we used the same strategy proposed in Petrou and García Sevilla (2006): a pixel belongs to a neighbourhood of radius d if its distance to the central pixel is in the range $[d - 1/2, d + 1/2)$. With this setting, the number of pixels forming the neighbourhood depends on d, and asymptotically approaches $[\pi d]$, where $[\cdot]$ indicates 'the nearest integer of'. The resulting neighbourhoods for $d = 1, 2, 3$ are shown in Fig. 3. In the remainder, we use the subscript '$\circ$' to refer to this variation.

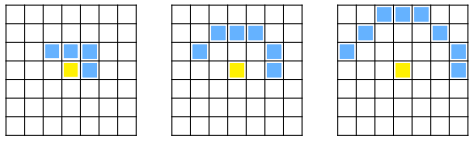

Fig. 3. Type of neighbourhood: digital circles.

2.2. Extraction of rotation dependent features

Consider a generic grey-scale image $\mathbf{I}$. From any neighbourhood containing N peripheral pixels we get N angular dependent co-occurrence matrices, each corresponding to the direction defined by pixels 0 and j (see Fig. 1). Now let $\mathbf{M}_j$ indicate any such co-occurrence matrix and $f_j^{(k)}$ a generic parameter extracted from it, with $1 \leq k \leq K$, being K the total number of parameters (let these, for instance, be *contrast*, *correlation*, *energy*, etc. – see Sec. 3.2). Ideally, as the input image rotates by angular steps multiple of the angle formed by two adjacent peripheral pixels, each k-th vector $\mathbf{f}^{(k)} = \left[f_1^{(k)}, \ldots, f_N^{(k)}\right]$ circularly shifts by one or more positions. In practice some intrinsic limitations of the digital domain do not let this condition hold perfectly: on the one hand, the length of the displacement vector varies as the direction changes; on the other hand, the angular interval is not uniform. It is nonetheless reasonable to assume that the circular shift condition holds, at least with good approximation. To obtain rotation invariant features we need to transform the $\mathbf{f}^{(k)}$ through functions that are invariant under a cyclic substitution. In the following subsection we discuss four different solutions.

2.3. Conversion from rotation dependent into rotation independent features

Let us start by recalling some basic functions that will serve as a basis for defining rotation independent features. Let $\mathbf{x} = [x_1, \ldots, x_N]$ denote a vector of real numbers; the functions are:

Average

$$\bar{\mathbf{x}} = \frac{1}{N} \sum_{j=1}^{N} x_j \tag{1}$$

Range

$$\Delta\mathbf{x} = \max\left(x_j\right) - \min\left(x_j\right); \qquad j \in \{1, \ldots, x_j, \ldots, N\} \tag{2}$$

Mean absolute deviation:

$$\delta\mathbf{x} = \frac{1}{N} \sum_{j=1}^{N} |x_j - \bar{\mathbf{x}}| \tag{3}$$

Discrete Fourier transform:

$$\hat{\mathbf{x}} = [\hat{x}_1, \ldots, \hat{x}_m, \ldots, \hat{x}_N]; \quad \hat{x}_m = \sum_{j=1}^{N} x_j e^{-i\frac{2\pi(m-1)(j-1)}{N}} \tag{4}$$

where $i = \sqrt{-1}$.

The above functions can be viewed as basic blocks for computing rotation invariant features. The use of average and range was advocated by Haralick himself in his seminal work (Haralick et al., 1973). Curiously, it seems that related literature has taken in only the first but, for reasons unknown to us, largely overlooked the second. The common treatment is in fact based on the average alone, as in Haralick et al. (1973); Mendoza et al. (2007); Xian (2010); Alam and Faruqui (2011); Bino Sebastian et al. (2012). In our experiments we considered both average alone and in concatenation with range. Finally, the discrete Fourier transform is another means to obtain rotation invariant features, for the moduli of the transformed coefficients $\hat{x}_m$ (Eq. 4) are invariant to any circular shift of the input vector $\mathbf{x}$ (see, for instance, Bianconi et al. 2009). To the best of our knowledge the mean absolute deviation and the absolute values of DFT coefficients have been not experimented before with co-occurrence matrices.

In summary we considered the following four variations to convert the rotation dependent features $\mathbf{f}^{(k)}$ into rotation independent ones: 1) the average of $\mathbf{f}^{(k)}$; 2) a concatenation of the average and range of $\mathbf{f}^{(k)}$; 3) a concatenation of the average, range and mean absolute deviation of $\mathbf{f}^{(k)}$; and 4) the absolute values of the DFT-transformed coefficients of $\mathbf{f}^{(k)}$. By combining these four feature normalization procedures with the two types of neighbourhoods described in Sec. 2.1 we get the eight treatments summarized in Tab. 2, where symbol '||' stands for 'concatenation'.

Table 2. Summary of the treatments studied in the experiments.

Neighbourhood type	Feature normalization	Symbol	No. of features
Four directions	Average	$\bar{\mathbf{f}}_*$	K
	Average + range	$\bar{\mathbf{f}}_* \| \Delta\mathbf{f}_*$	$2K$
	Average + range + mean absolute deviation	$\bar{\mathbf{f}}_* \| \Delta\mathbf{f}_* \| \delta\mathbf{f}_*$	$3K$
	Discrete Fourier transform	$\hat{\mathbf{f}}_*$	$4K$
Digital circles	Average	$\bar{\mathbf{f}}_\circ$	K
	Average + range	$\bar{\mathbf{f}}_\circ \| \Delta\mathbf{f}_\circ$	$2K$
	Average + range + mean absolute deviation	$\bar{\mathbf{f}}_\circ \| \Delta\mathbf{f}_\circ \| \delta\mathbf{f}_\circ$	$3K$
	Discrete Fourier transform	$\hat{\mathbf{f}}_\circ$	$\approx [\pi d]\, K$

3. Experiments

We performed a set of image classification experiments to evaluate accuracy and robustness against rotation of the descriptors presented in the preceding section. Datasets and procedure used in the experiments are detailed in the following subsections.

3.1. Datasets

The four datasets used in the experiments and summarized in Secs. 3.1.1 – 3.1.4 contain hardware-rotated images of planar or almost planar surfaces showing approximately stationary textures. These inclusion criteria are supported by the following considerations: first, the need to avoid software-rotated images, since rotation by software may modify the image microstructure and lead to misleading results (Fernández et al., 2011); second, images from planar or almost planar surfaces are less prone to artifacts resulting from shadows caused by surface roughness; finally, the use of stationary textures minimizes the sampling error related to the creation of train and validation sets (see Sec. 3.3).

3.1.1. Brodatz

This dataset contains the following 13 texture classes from the Brodatz's album: D9, D12, D15 D16, D19, D24, D29, D38, D68, D84, D92, D94 and D112. Hardware-rotated digital images have been captured in our lab directly from the original book (Brodatz, 1966) through the imaging system described in Bianconi et al. 2013. The apparatus permits to acquire rotated images at angular steps of 10°, a feature that we used to take images at the following rotation angles: 0°, 10°, 20°, 30°, 40°, 50°, 60°, 70°, 80° and 90°. The original images has been subdivided into 4×4 non-overlapping samples of resolution 205×205 pixels each, resulting in 16 samples per class.

3.1.2. MondialMarmi

MondialMarmi is a free image database of granite tiles for colour and texture analysis. The current version (1.1) includes 12 granite classes: Acquamarina, Azul Capixaba, Azul Platino, Bianco Cristal, Bianco Sardo, Giallo Napoletano, Giallo Ornamentale, Giallo Santa Cecilia, Giallo Veneziano, Rosa Beta, Rosa Porrio A, Rosa Porrio B. The dataset features hardware-rotated images taken at nine different rotation angles: 0°, 5°, 10°, 15°, 30°, 45°, 60°, 75° and 90° (for a detailed description of the acquisition procedure see Fernández et al. 2011). There are four images for each class, every one image representing a different tile. The original images have been subdivided into four non-overlapping samples generating a dataset of 16 samples per class of resolution 272×272.

3.1.3. Outex

Outex is a well-established texture database for evaluating texture classification and segmentation algorithms. The suite contains hardware-rotated images taken at the same rotation angles as MondialMarmi. Herein we considered a selection of 45 texture classes: canvas{005, 021}; carpet005; granite{001, 003, 004, 005, 006, 007, 008, 009, 010}; paper006; plastic{001, 002, 003, 004, 005, 009, 016, 017, 018, 019, 020, 021, 022, 023, 024, 025, 026, 027, 028, 029, 030, 031, 032, 033, 034, 035, 036, 038, 040, 041}; wood{006, 008}. We refer to this subset of the Outex database as OUTEX_00045. This selection satisfies, to a great extent, the inclusion/exclusion criteria stated at the beginning of the section (i.e.: flatness of the original surface and stationariness). Following the instructions available in OuTeX, we subdivided the original images into 20 non-overlapping subsamples of resolution 128×128 each.

3.1.4. Vectorial

This dataset is composed of 20 artificial texture classes derived from vectorial images (Fig. 4). These have been downloaded from a free-access repository (All-free-download). Unlike raster images, vectorial images can be software-rotated without introducing any artifacts. Therefore, images rotated this way can be considered as if they were hardware-rotated. We constructed this dataset in the following way: first, we resized the original images to 20cm × 20cm; second, we rotated each image by the same angles used in the Brodatz dataset through a free editor for vector graphics (Inkscape); third, we rasterized the resulting images at a resolution of 300dpi; finally, we retained a centred square of each raster image and subdivided it into 16 non-overlapping sub-images of dimension 225×225 each.

3.2. Features

For each direction resulting from a certain neighbourhood type and a given distance d, we computed five features among those most frequently used in the literature, namely: *contrast*,

Table 3. Overall results.

Texture descriptor	No. of features	Dataset			
		Brodatz	MondialMarmi	OUTEX_00045	Vectorial
$\bar{\mathbf{f}}'_{1^*}$	5	77.7 ± 2.6	80.8 ± 2.1	56.4 ± 1.2	77.7 ± 3.2
$\bar{\mathbf{f}}_{1^*}$	5	76.6 ± 2.2	80.5 ± 2.1	56.4 ± 1.2	77.6 ± 3.2
$\bar{\mathbf{f}}_{1^*} \Vert \Delta\mathbf{f}_{1^*}$	10	78.8 ± 13.0	86.5 ± 2.4	59.4 ± 4.2	80.9 ± 3.6
$\bar{\mathbf{f}}_{1^*} \Vert \Delta\mathbf{f}_{1^*} \Vert \delta\mathbf{f}_{1^*}$	15	78.6 ± 13.0	86.7 ± 2.6	59.8 ± 4.2	80.9 ± 3.6
$\hat{\mathbf{f}}_{1^*}$	20	82.4 ± 8.8	84.9 ± 2.2	59.7 ± 3.4	80.4 ± 3.3
$\bar{\mathbf{f}}'_{1^\circ}$	5	77.7 ± 2.6	80.8 ± 2.1	56.4 ± 1.2	77.7 ± 3.2
$\bar{\mathbf{f}}_{1^\circ}$	5	76.6 ± 2.2	80.5 ± 2.1	56.4 ± 1.2	77.6 ± 3.2
$\bar{\mathbf{f}}_{1^\circ} \Vert \Delta\mathbf{f}_{1^\circ}$	10	78.8 ± 13.0	86.5 ± 2.4	59.4 ± 4.2	80.9 ± 3.6
$\bar{\mathbf{f}}_{1^\circ} \Vert \Delta\mathbf{f}_{1^\circ} \Vert \delta\mathbf{f}_{1^\circ}$	15	78.6 ± 13.0	86.7 ± 2.6	59.8 ± 4.2	80.9 ± 3.6
$\hat{\mathbf{f}}_{1^\circ}$	20	82.4 ± 8.8	84.9 ± 2.2	59.7 ± 3.4	80.4 ± 3.3
$\bar{\mathbf{f}}'_{2^*}$	5	82.0 ± 2.7	81.9 ± 2.6	55.9 ± 0.6	81.0 ± 2.4
$\bar{\mathbf{f}}_{2^*}$	5	80.6 ± 2.9	81.6 ± 2.7	56.3 ± 0.7	80.8 ± 2.4
$\bar{\mathbf{f}}_{2^*} \Vert \Delta\mathbf{f}_{2^*}$	10	84.2 ± 8.8	88.8 ± 2.3	59.9 ± 3.6	83.1 ± 3.3
$\bar{\mathbf{f}}_{2^*} \Vert \Delta\mathbf{f}_{2^*} \Vert \delta\mathbf{f}_{2^*}$	15	83.7 ± 9.2	89.3 ± 2.3	**60.4** ± 3.6	83.2 ± 3.4
$\hat{\mathbf{f}}_{2^*}$	20	87.5 ± 6.3	88.4 ± 2.3	**60.4** ± 2.5	83.5 ± 2.8
$\bar{\mathbf{f}}'_{2^\circ}$	5	83.2 ± 1.9	82.1 ± 2.7	56.2 ± 0.8	80.6 ± 2.3
$\bar{\mathbf{f}}_{2^\circ}$	5	81.0 ± 1.7	81.9 ± 2.7	56.6 ± 0.7	80.5 ± 2.3
$\bar{\mathbf{f}}_{2^\circ} \Vert \Delta\mathbf{f}_{2^\circ}$	10	90.5 ± 6.1	87.8 ± 3.1	57.1 ± 6.5	83.0 ± 3.0
$\bar{\mathbf{f}}_{2^\circ} \Vert \Delta\mathbf{f}_{2^\circ} \Vert \delta\mathbf{f}_{2^\circ}$	15	90.8 ± 6.1	88.0 ± 3.1	57.3 ± 6.5	82.9 ± 3.1
$\hat{\mathbf{f}}_{2^\circ}$	30	93.3 ± 3.0	87.4 ± 2.4	59.5 ± 3.0	83.4 ± 2.3
$\bar{\mathbf{f}}'_{3^*}$	5	82.2 ± 2.6	79.9 ± 2.0	52.0 ± 0.7	82.4 ± 1.7
$\bar{\mathbf{f}}_{3^*}$	5	81.0 ± 2.3	79.4 ± 2.1	52.2 ± 0.6	82.2 ± 1.8
$\bar{\mathbf{f}}_{3^*} \Vert \Delta\mathbf{f}_{3^*}$	10	85.8 ± 7.3	88.3 ± 1.5	56.6 ± 1.7	84.3 ± 3.1
$\bar{\mathbf{f}}_{3^*} \Vert \Delta\mathbf{f}_{3^*} \Vert \delta\mathbf{f}_{3^*}$	10	85.6 ± 7.4	89.0 ± 1.6	57.6 ± 1.8	84.3 ± 3.2
$\hat{\mathbf{f}}_{3^*}$	15	87.7 ± 6.7	88.8 ± 1.8	58.1 ± 1.4	84.5 ± 2.5
$\bar{\mathbf{f}}'_{3^\circ}$	5	84.2 ± 1.1	80.9 ± 2.3	54.4 ± 0.9	82.2 ± 1.9
$\bar{\mathbf{f}}_{3^\circ}$	5	81.9 ± 1.4	80.7 ± 2.5	54.8 ± 1.0	82.0 ± 1.9
$\bar{\mathbf{f}}_{3^\circ} \Vert \Delta\mathbf{f}_{3^\circ}$	10	94.7 ± 2.6	89.6 ± 1.9	58.6 ± 1.5	84.7 ± 2.0
$\bar{\mathbf{f}}_{3^\circ} \Vert \Delta\mathbf{f}_{3^\circ} \Vert \delta\mathbf{f}_{3^\circ}$	15	94.8 ± 2.5	**89.7** ± 1.9	58.6 ± 1.7	84.7 ± 2.0
$\hat{\mathbf{f}}_{3^\circ}$	40	95.1 ± 1.2	88.6 ± 1.9	58.7 ± 1.5	85.0 ± 1.7
$\bar{\mathbf{f}}'_{4^*}$	5	79.8 ± 2.5	78.3 ± 2.2	49.7 ± 1.2	82.3 ± 1.9
$\bar{\mathbf{f}}_{4^*}$	5	78.3 ± 2.9	78.0 ± 2.3	49.8 ± 1.3	82.1 ± 2.0
$\bar{\mathbf{f}}_{4^*} \Vert \Delta\mathbf{f}_{4^*}$	10	79.8 ± 12.4	86.9 ± 2.2	54.9 ± 1.0	84.5 ± 3.2
$\bar{\mathbf{f}}_{4^*} \Vert \Delta\mathbf{f}_{4^*} \Vert \delta\mathbf{f}_{4^*}$	15	79.6 ± 12.5	87.5 ± 2.2	55.7 ± 1.2	84.5 ± 3.3
$\hat{\mathbf{f}}_{4^*}$	20	82.9 ± 11.1	88.6 ± 1.6	56.6 ± 1.3	84.4 ± 2.6
$\bar{\mathbf{f}}'_{4^\circ}$	5	83.5 ± 1.4	79.4 ± 2.0	52.2 ± 0.7	82.9 ± 1.7
$\bar{\mathbf{f}}_{4^\circ}$	5	82.3 ± 1.5	78.9 ± 2.0	52.1 ± 0.8	82.7 ± 1.8
$\bar{\mathbf{f}}_{4^\circ} \Vert \Delta\mathbf{f}_{4^\circ}$	10	95.2 ± 0.7	87.9 ± 1.0	59.9 ± 0.9	**86.0** ± 1.6
$\bar{\mathbf{f}}_{4^\circ} \Vert \Delta\mathbf{f}_{4^\circ} \Vert \delta\mathbf{f}_{4^\circ}$	15	**95.4** ± 0.7	88.2 ± 1.1	60.0 ± 0.9	**86.0** ± 1.7
$\hat{\mathbf{f}}_{4^\circ}$	80	94.6 ± 1.0	88.8 ± 1.5	59.9 ± 1.0	85.6 ± 1.4

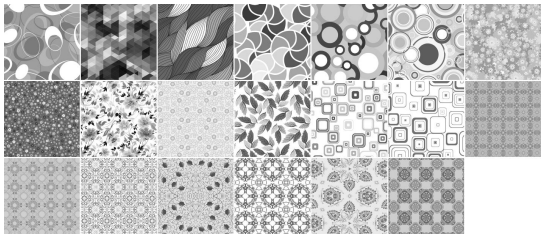

Fig. 4. Vectorial dataset (20 classes).

correlation, *energy*, *entropy* and *homogeneity*. The implementation adopted herein guarantees that all the features have range [0, 1]. For this reason our implementation can be slightly different from other formulations in which features are normalized in a different way. In formulas we have:

Contrast:

$$f^{(1)}(\mathbf{M}) = \frac{1}{(G-1)^2} \sum_{u=0}^{G-1} \sum_{v=0}^{G-1} |u-v|^2 \mathbf{M}(u,v) \tag{5}$$

Correlation:

$$f^{(2)}(\mathbf{M}) = \frac{\sum_{u=0}^{G-1} \sum_{v=0}^{G-1} (u-\mu_u)(v-\mu_v)\mathbf{M}(u,v)}{2\sigma_u\sigma_v} + \frac{1}{2} \tag{6}$$

Energy

$$f^{(3)}(\mathbf{M}) = \sum_{u=0}^{G-1} \sum_{v=0}^{G-1} \mathbf{M}(u,v)^2 \tag{7}$$

Entropy

$$f^{(4)}(\mathbf{M}) = -\frac{1}{2\log_2(G)} \sum_{u=0}^{G-1} \sum_{v=0}^{G-1} \mathbf{M}(u,v) \log_2 [\mathbf{M}(u,v)] \tag{8}$$

Homogeneity

$$f^{(5)}(\mathbf{M}) = \sum_{u=0}^{G-1}\sum_{v=0}^{G-1}\frac{\mathbf{M}(u,v)}{1+|u-v|} \quad (9)$$

where G represents the number of grey levels; u,v the coordinates of the co-occurrence matrix; μ_u, μ_v, σ_u and σ_v the means and the standard deviations of the marginal distributions associated with the co-occurrence matrix.

3.3. Estimation of classification accuracy

We estimated the classification accuracy through a supervised classification task based on the nearest-neighbour (1-NN) rule with L_1 distance. The absence of tuning parameters, as well as the ease of implementation and other desirable asymptotic properties, support and explain the wide use of this method for comparative purposes (see for example Varma and Zisserman 2009; Guo et al. 2010; Crosier and Griffin 2010; Kandaswamy et al. 2011). For each dataset we generate a predefined number of classification problems (100 in our experiments); in each problem we split the dataset into two non-overlapping subsets, one for training and the other for validation, with the constraint that, for each class, half of the samples are used for training and the other half for validation. To study the effects of image rotation, we always train the classifier with unrotated images and test it using images rotated by θ degrees, where θ is one of the rotation angles available in the dataset.

3.4. Calibration against other methods

For comparison purposes we calibrated the results against the following rotation-invariant texture descriptors: circular and radial co-occurrence features as proposed by Barrera et al. (2012); wavelet co-occurrence features (WCF) as proposed by Mukane et al. (2011); rotation-invariant local binary patterns (versions $\mathrm{LBP}^{ri}_{8,1}$ and $\mathrm{LBP}^{ri}_{16,2}$). We briefly outline the first two methods in the following subsections; for details about LBP the interested reader is referred to the work Ojala et al. (2002). The accuracy of each method is reported in Tab. 4.

3.4.1. Circular and radial co-occurrence features

The method described by Barrera et al. (2012) is based on a combination of 'circular' and 'radial' features. Circular features encode the co-occurrence of the mean grey-level of two concentric circles; radial features the co-occurrence of the mean grey-level of adjacent radial lines separated by uniform angular spacing. From the resulting co-occurrence matrices suitable statistics are extracted from each of the two groups; then, for any pair of them, a new feature is computed as the root mean square value of each pair. In our implementation we maintained the settings recommended in the cited reference, namely an internal and external radius value $r_i = 2$ and $r_e = 2$, respectively, for the circular features, and $N = 8$ directions with radius $r = 5$ for the radial features. For fair play we used the same five statistics described in Sec. 3.2. As a result, the number of features resulting from this implementation is five.

3.4.2. Wavelet co-occurrence features

This approach combines standard grey-level co-occurrence features computed not only from the original image, but also from each of the four sub-bands resulting from one-level wavelet decomposition of the original image. As suggested in the original reference (Mukane et al., 2011), we perform image decomposition through Daubechies' 'db4' wavelet. In our implementation we compute co-occurrences using the four standard directions (i.e. 0°, 45°, 90° and 135°) and distance $d = \{1, \ldots, 4\}$. Then the five statistics described in Sec. 3.2 are extracted from each co-occurrence matrix. For rotation-invariance we followed the approach proposed by Zhu et al. (2011), which consists of extracting the mean and standard deviation of each statistic over the different directions. As a result we get ten features from the original image and for each of the four one-level sub-bands, giving a total number of 50 features as a result.

3.5. Reproducible research

For reproducible research purposes, all the data required to replicate the experiments (i.e.: source code, images and subdivisions into train and validation sets) are available online (GLCM)[2].

4. Results and discussion

The overall results of the classification experiment are summarized in Tab. 3. For each descriptor and dataset the table reports the results in the form $\mu \pm \sigma$, where μ and σ are, respectively, the mean and standard deviation of the classification accuracy (in %) over the different rotation angles. Numeric subscripts indicate the distance (d) at which the treatments have been computed. The complete results for each rotation angle are available in the companion website (GLCM). To draw meaningful conclusions we pairwise compared the variations from the two standpoints of accuracy and robustness, as detailed below.

4.1. Pairwise comparisons: accuracy

To pairwise compare accuracies within the same dataset we adopted the following rule: a treatment T_a outperforms a treatment T_b if the accuracy of T_a is, for each angle, significantly superior to that of T_b. This is checked through Wilcoxon's signed rank sum test for equal medians with $\alpha = 0.05$. The test takes as input the two vectors containing the 100 accuracy values produced by each descriptor in the 100 classification problems (see Sec. 3.3).

To evaluate the response to each design factor we proceeded in the following way: let V_a and V_b be two variations of a design factor (let for instance V_a indicate 'neighbourhood type/four directions' and V_b 'neighbourhood type/digital circles'). Each time a treatment based on V_a outperforms one based on V_b we assign a win to V_a and a draw to V_b (and vice-versa); otherwise we assign a draw. Wins, draws and losses are summed up for

[2]To access the page: user = `co-occurrence`, password = `rotation`

Table 4. Comparison with other methods.

Texture descriptor	No. of features	Dataset			
		Brodatz	MondialMarmi	OUTEX_00045	Vectorial
Circular + radial co-occurrence features	5	60.9 ± 12.6	75.8 ± 2.7	35.0 ± 2.3	69.8 ± 2.3
$WCF_{d=1}$	50	77.4 ± 11.8	78.9 ± 0.9	70.7 ± 7.5	80.0 ± 2.9
$WCF_{d=2}$	50	82.0 ± 10.6	80.3 ± 1.9	73.9 ± 7.2	79.9 ± 2.9
$WCF_{d=3}$	50	79.7 ± 13.7	79.8 ± 2.7	71.2 ± 8.3	80.0 ± 2.9
$WCF_{d=4}$	50	73.1 ± 16.6	80.0 ± 3.0	70.4 ± 7.7	**80.1** ± 2.8
$LBP^{ri}_{8,1}$	36	83.0 ± 7.7	83.6 ± 2.4	73.7 ± 1.2	70.4 ± 5.8
$LBP^{ri}_{16,2}$	4116	**95.0** ± 1.9	**88.5** ± 1.5	**78.8** ± 1.6	75.3 ± 5.5

all distances and datasets (Tab. 5). As for neighbourhood type, note that the comparison excludes the case $d = 1$, a situation in which the two variations generate the same neighbourhood (see Figs. 2 and 3). In addition, the table reports the main effect (averaged over the the four datasets) that we get when we switch from V_a to V_b. The figures are in percentage points (pp).

Table 5. Pairwise comparisons: accuracy.

Variation V_a	Variation V_b	Wins V_a	Draws	Wins V_b	Main effect (pp) $V_a \rightarrow V_b$
$\mathbf{f}_*$	$\mathbf{f}_\circ$	27	113	52	+1.8
$\bar{\mathbf{f}}$	$\bar{\mathbf{f}}\|\|\Delta\mathbf{f}$	0	26	38	+5.1
$\bar{\mathbf{f}}$	$\bar{\mathbf{f}}\|\|\Delta\mathbf{f}\|\|\delta\mathbf{f}$	0	27	37	+5.3
$\bar{\mathbf{f}}$	$\hat{\mathbf{f}}$	0	18	46	+5.8
$\bar{\mathbf{f}}\|\|\Delta\mathbf{f}$	$\bar{\mathbf{f}}\|\|\Delta\mathbf{f}\|\|\delta\mathbf{f}$	2	58	4	+0.2
$\bar{\mathbf{f}}\|\|\Delta\mathbf{f}$	$\hat{\mathbf{f}}$	1	58	5	+0.7
$\bar{\mathbf{f}}\|\|\Delta\mathbf{f}\|\|\delta\mathbf{f}$	$\hat{\mathbf{f}}$	1	62	1	+0.5

4.2. Pairwise comparison: robustness

With the term 'robustness' we mean the capability of a descriptor to maintain the same accuracy when the relative orientation between train and test images changes. A reasonable estimate of this property is the standard deviation of the classification accuracy over the different rotation angles (see Tab. 3). A high variability means low robustness and vice-versa. For a given dataset we say that a treatment T_a is more robust than a treatment T_b if the standard deviation of the accuracy of T_a over the rotation angles featured by the dataset is significantly lower than that of T_b. To state whether there is significant difference between two standard deviations, we used the F-test ($\alpha = 0.05$). Then we proceeded like in the previous section: each time a treatment based on V_a is more robust than one based on V_b we assign a win to V_a and a draw to V_b (and vice-versa); otherwise we assign a draw. Again, wins, draws and losses are summed up for all distances and datasets (Tab. 6). Note that a negative main effect indicates an improvement in robustness.

Evaluation of robustness can be considered a second, 'finer' step to compare descriptors and treatments. It can be used to further discriminate among methods that possess similar accuracy. For this reason we only reported, in Tab. 6, robustness comparison between the methods that emerged as the most

Table 6. Pairwise comparisons: robustness.

Treatment V_a	Treatment V_b	Wins V_a	Draws	Wins V_b	Main effect (pp) $V_a \rightarrow V_b$
$\mathbf{f}_*$	$\mathbf{f}_\circ$	16	115	61	-1.0
$\bar{\mathbf{f}}\|\|\Delta\mathbf{f}$	$\bar{\mathbf{f}}\|\|\Delta\mathbf{f}\|\|\delta\mathbf{f}$	4	56	4	+0.1
$\bar{\mathbf{f}}\|\|\Delta\mathbf{f}$	$\hat{\mathbf{f}}$	2	53	9	-0.9
$\bar{\mathbf{f}}\|\|\Delta\mathbf{f}\|\|\delta\mathbf{f}$	$\hat{\mathbf{f}}$	2	53	9	-0.9

prominent from the pairwise comparisons of accuracies. Finally, the plots in Fig. 5 summarize the main effects on accuracy and robustness as functions of d.

4.3. Discussion

The average accuracy values reported in Tab. 3 confirm the trend recently brought to light by Kandaswamy et al. (2011), namely that co-occurrence features are fairly accurate in classification problems involving relatively few classes. This is true, in our case, with datasets Brodatz, MondialMarmi and Vectorial (respectively composed of 13, 12 and 20 classes) for which we easily get over 80% accuracy. With more classes the figures drop significantly: with OUTEX_00045 (45 classes) we barely attain 60%. Comparison with alternative rotationally-invariant co-occurrence features and other rotation-invariant descriptors (see Tab. 4) shows that the formulations discussed in this paper outperform the other methods in three datasets out of four.

Beyond absolute accuracy, it is however more relevant to this work to analyse the relative effects of the treatments presented in the preceding section.

4.3.1. Effect of the type of neighbourhood

As one would expect, switching from the original formulation (four directions) to digital circles brings beneficial effects both in terms of accuracy and robustness. As for accuracy, digital circles are almost twice more likely (52/27) to outperform the original formulation, with an overall main effect of +1.8 percentage points. The difference is more marked for robustness: in this case digital circles are approximately four times more likely (61/16) to produce more robust results than the original formulation. Fig. 5 shows that the better performance of digital circles compared with the original formulation is more noticeable as d increases. This result makes sense too, since the

original formulation always samples the input image at fixed angular intervals of 45°, no matter the value of d, whereas digital circles produce a finer angular sampling of the input image when d increases.

4.3.2. *Effect of the type of conversion from rotation dependent into rotation invariant features*

The main and clearest result here is that the most common approach for obtaining rotation invariant features (i.e.: taking the average of the rotation dependent features) is, by far, the worst. In no case the accuracy of this method is superior to any of the other normalization methods considered in our experiments (see Tab. 5). The main effect of switching from average alone to any other method is also conspicuous, ranging from +5.1 to +5.8 percentage points. It is important to emphasize, at this point, that some implementations proposed in literature use the average of the co-occurrence matrices themselves, instead of the average of the features. In our experiments we found no significant difference between the two methods (for the complete results see GLCM), therefore, for the sake of homogeneity with the others, we retained the second. The remaining three approaches exhibit similar results, in terms of accuracy, with a little gap in favour of the discrete Fourier transform. This effect is more marked when we speak of robustness (Tab. 6): DFT is approximately four times (9/2) more likely to produce more robust results than either average+range or average+range+mean average deviation, with an average decrease in σ of 0.9 percentage points. Fig. 5 shows no clear trend as a function of d.

5. Conclusions

Co-occurrence matrices are among the most used image descriptors, with applications covering almost any area of image analysis. They are also one of the oldest – they first appeared on the scene forty years ago. In this study we investigated whether it was possible to improve robustness and accuracy of the method for rotation invariant image classification. The results of our study are in the affirmative: there is room for improvement.

As for the type of neighbourhood, digital circles proved appreciably superior to the original formulation (four directions). As one would expect, the beneficial effect is more noticeable as d increases. It is therefore recommendable to adopt this setting for any $d > 1$. Regarding the conversion into rotation invariant features, the use of the average alone is significantly inferior to any other methods studied here, and should therefore be avoided. Any of the three other methods studied here (i.e.: average + range, average + range + mean absolute deviation or discrete Fourier transform) works better. Of these, the DFT (to the best of our knowledge not studied before in this context) produced globally better results both in terms of accuracy and robustness, though at the cost of a higher number of features (see. Tab. 2). Comparison with other methods suggest that, when properly implemented, co-occurrence features may even perform better than well reputed descriptors such as LBP.

In summary, for practical applications requiring invariance against rotation, we recommend using digital circles instead of the four-direction neighbourhood. We advise DFT normalization if the number of features is not critical; otherwise, average + range or average + range + mean absolute deviation can be used instead. We discourage the use of the most popular approach (i.e. average over four directions), since any of the other treatments considered in this study performs better.

References

Alam, F., Faruqui, R., 2011. Optimized calculations of Haralick texture features. European Journal of Scientific Research 50, 543–553.

All-free-download, 2012. All-free-download. `http://all-free-download.com`.

Arebey, M., Hannan, M., Begum, R., Basri, H., 2012. Solid waste bin level detection using gray level co-occurrence matrix feature extraction approach. Journal of Environmental Management 104, 9–18.

Azzabou, N., Paragios, N., Guichard, F., 2010. Image reconstruction using particle filters and multiple hypotheses testing. IEEE Transactions on Image Processing 19, 1181–1190.

Barrera, M., Andrade, M., Kim, H., 2012. Texture-based fuzzy system for rotation-invariant classification of aerial orthoimage regions, in: Proceedings of the 4th International Conference on Geographic Object-Based Image Analysis (GEOBIA 2012), Rio de Janeiro, Brazil. pp. 58–62.

Barwad, A., Dey, P., Susheilia, S., 2012. Artificial neural network in diagnosis of metastatic carcinoma in effusion cytology. Cytometry Part B (Clinical Cytometry) 82, 107–111.

Ben Salem, Y., Nasri, S., 2010. Automatic recognition of woven fabrics based on texture and using SVM. Signal, Image and Video Processing 4, 429–434.

Berthelot, A., Solberg, A., Gelius, L., 2013. Texture attributes for detection of salt. Journal of Applied Geophysics 88, 52–69.

Bianconi, F., Fernández, A., González, E., Armesto, J., 2009. Robust colour texture features based on ranklets and discrete Fourier transform. Journal of Electronic Imaging 18, 043012–1–8. doi:`10.1117/1.3273946`.

Bianconi, F., Fernández, A., González, E., Saetta, S.A., 2013. Performance analysis of colour descriptors for parquet sorting. Expert Systems with Applications 40, 1636–1644. doi:`10.1016/j.eswa.2012.09.007`.

Bino Sebastian, V., Unnikrishnan, A., Balakrishnan, K., 2012. Gray level co-occurrence matrices: Generalisation and some new features. International Journal of Computer Science, Engineering and Information Technology 2, 151–157.

Brodatz, P., 1966. Textures: a photographic album for artists and designers. Dover Publications, New York, USA.

Clausi, D.A., Deng, H., 2005. Design-based texture feature fusion using Gabor filters and co-occurrence probabilities. IEEE Transactions on Image Processing 14, 925–936.

Cochran, W., Cox, G., 1957. Experimental designs. Second edition. John Wiley & Sons, Inc., New York.

Crosier, M., Griffin, L.D., 2010. Using basic image features for texture classification. International Journal of Computer Vision 88, 447–460.

Davis, L.S., Johns, S.A., Aggarwal, J.K., 1979. Texture analysis using generalized co-occurrence matrices. IEEE Transactions on Pattern Analysis and Machine Intelligence PAMI-1, 251–259.

Dutta, S., Datta, A., Chakladar, N.D., Pal, S., Mukhopadhyay, S., Sen, R., 2012. Detection of tool condition from the turned surface images using an accurate grey level co-occurrence technique. Precision Engineering 36, 458–466.

Fernández, A., Ghita, O., González, E., Bianconi, F., Whelan, P.F., 2011. Evaluation of robustness against rotation of LBP, CCR and ILBP features in granite texture classification. Machine Vision and Applications 22, 913–926. doi:`10.1007/s00138-010-0253-4`.

GLCM, 2013. GLCM. Code, data and results related to this paper. Available online at `http://webs.uvigo.es/antfdez/downloads.html/`.

Gómez, W., Pereira, W., Infantosi, A., 2012. Analysis of co-occurrence texture statistics as a function of gray-level quantization for classifying breast ultrasound. IEEE Transactions on Medical Imaging 31, 1889–1899.

González, E., Fernández, A., Bianconi, F., 2014. General framework for rotation invariant texture classification through co-occurrence of patterns. Journal of Mathematical Imaging and Vision doi:`10.1007/s10851-014-0500-9`. in press.

Guo, Z., Zhang, L., Zhang, D., 2010. A completed modeling of local binary pattern operator for texture classification. IEEE Transactions on Image Processing 19, 1657–1663.

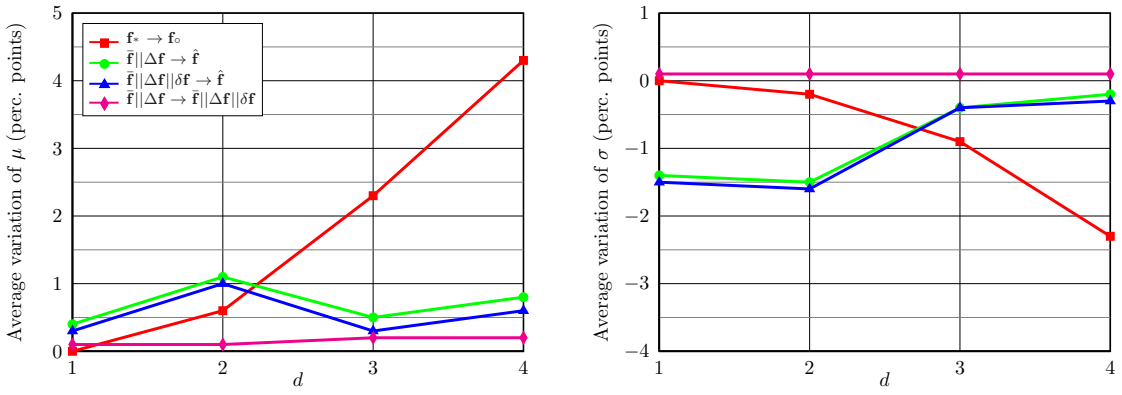

Fig. 5. Main effects of the variations on accuracy (left) and robustness (right) as functions of d.

Haralick, R.M., Shanmugam, K., Dinstein, I., 1973. Textural features for image classification. IEEE Transactions on Systems, Man, and Cybernetics 3, 610–621.

Hu, W., Li, H., Wang, C., Gou, S., Fu, L., 2012. Characterization of collagen fibers by means of texture analysis of second harmonic generation images using orientation-dependent gray level co-occurrence matrix method. Journal of Biomedical Optics 17, 026007–1–8.

Inkscape, 2012. Inkscape: open source scalable vector graphics editor. `http://www.inkscape.org`.

Kandaswamy, U., Adjeroh, D., Lee, M., 2005. Efficient texture analysis of SAR imagery. IEEE Transactions on Geoscience and Remote Sensing 43, 2075–2083.

Kandaswamy, U., Schuckers, S.A., Adjeroh, D., 2011. Comparison of texture analysis schemes under nonideal conditions. IEEE Transactions on Image Processing 20, 2260–2275.

Linder, N., Konsti, J., Turkki, R., Rahtu, E., Lundin, M., Nordling, S., Haglund, C., Ahonen, T., Pietikäinen, M., Lundin, J., 2012. Identification of tumor epithelium and stroma in tissue microarrays using texture analysis. Diagnostic Pathology 7, 1–11.

Manivannan, K., Aggarwal, P., Devabhaktuni, V., Kumar, A., Nims, D., Bhattacharya, P., 2012. Particulate matter characterization by gray level co-occurrence matrix based support vector machines. Journal of Hazardous Materials 223-224, 94–103.

Masetti, G., Calder, B., 2012. Remote identification of a shipwreck site from MBES backscatter. Journal of Environmental Management 111, 44–52.

Mendoza, F., Dejmek, P., Aguilera, J.M., 2007. Colour and image texture analysis in classification of commercial potato chips. Food Research International 40, 1146–1154.

Mukane, S., Bormane, D., Gengaje, S., 2011. Wavelet and co-occurrence matrix based rotation invariant features for texture image retrieval using fuzzy logic. International Journal of Computer Applications 24, 1–5.

Mukherjee, J., Das, P., Kumar, M.A., Chatterji, B., 2000. On approximating Euclidean metrics by digital distances in 2D and 3D. Pattern Recognition Letters 21, 573–582.

Ojala, T., Pietikäinen, M., Mäenpää, T., 2002. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. IEEE Transactions on Pattern Analysis and Machine Intelligence 24, 971–987.

OuTeX, 2002. OuTeX database. Available online at `http://www.outex.oulu.fi/`.

Palm, C., 2004. Color texture classification by integrative co-occurrence matrices. Pattern Recognition 37, 965–976.

Parekh, R., 2012. Using texture analysis for medical diagnosis. IEEE MultiMedia 19, 28–37.

Petrou, M., García Sevilla, P., 2006. Image Processing. Dealing with Texture. Wiley Interscience.

Ramamurthy, B., Chandran, K., 2012. Content based medical image retrieval with texture content using gray level co-occurrence matrix and k-means clustering algorithms. Journal of Computer Science 8, 1070–1076.

Song, Q., 2011. Illumination invariant texture classification with pattern co-occurrence matrix, in: Shen, G., Huang, X. (Eds.), Advanced Research on Computer Science and Information Engineering. Springer. volume 152 of *Communications in Computer and Information Science*, pp. 67–72.

Varma, M., Zisserman, A., 2009. A statistical approach to material classification using image patch exemplars. IEEE Transactions on Pattern Analysis and Machine Intelligence 31, 2032–2047.

Wu, Q., An, J., Lin, B., 2012. A texture segmentation algorithm based on PCA and global minimization active contour model for aerial insulator images. IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing 5, 1509–1518.

Xian, G., 2010. An identification method of malignant and benign liver tumors from ultrasonography based on GLCM texture features and fuzzy SVM. Expert Systems with Applications 37, 6737–6741.

Xie, X., Mirmehdi, M., 2008. A galaxy of texture features, in: Mirmehdi, M., Xie, X., Suri, J. (Eds.), Handbook of texture analysis. Imperial College Press, pp. 375–406.

Zhu, Z., Zhao, C., Hou, Y., Gao, H., 2011. Rotation-invariant texture image retrieval based on combined feature sets. International Journal of Digital Content Technology and its Applications 5, 287–292.