Variance Ranklets: orientation–selective rank features for contrast modulations

George Azzopardi
http://www.cs.rug.nl/~george

Fabrizio Smeraldi
http://www.dcs.qmul.ac.uk/~fabri

Department of Mathematics and
Computing Science
University of Groningen
Groningen, Netherlands

School of Electronic Engineering
and Computer Science
Queen Mary University of London
London, UK

Abstract

We introduce a novel type of orientation–selective rank features that are sensitive to contrast modulations (second–order stimuli). Variance Ranklets are designed in close analogy with the standard Ranklets, but use the Siegel–Tukey statistics for dispersion instead of the Wilcoxon statistics. Their response shows the same orientation selectivity pattern of Haar wavelets on second–order signals that are not detectable by linear filters. To the best of our knowledge, this is the first family of rank filters designed to detect orientation in variance modulations.

We validate our descriptors with an application to texture classification over a subset of the VisTex and Brodatz databases. The combination of standard (intensity) Ranklets with Variance Ranklets greatly improves on the performance of Ranklets alone. Comparison with other published results shows that state–of–the–art recognition rates can be achieved with a simple Nearest Neighbour classifier.

1 Introduction

Rank features are arguably the most robust kind of descriptors available to Computer Vision practitioners, being invariant to all monotonic transformations of image brightness, such as gamma correction and histogram equalisation. New, sophisticated kinds of descriptors have been developed in recent years, allowing applications to texture segmentation [1, 20] and object recognition [8] alongside the classical benchmark of wide-baseline stereo matching [4, 25].

Orientation selectivity has recently been introduced to rank features by Ranklets [22]. These descriptors are designed in close analogy with Haar wavelets; they are based on the Wilcoxon statistics and have a simple interpretation in terms of the pairwise comparison of pixel intensity values. Ranklets have been applied to face detection [23], point tracking [7], the processing of mammographic images [14, 17] and texture classification [15].

Orientation selectivity is one of the key features of widely used linear filters such as Haar wavelets, derivatives of Gaussians and Gabor filters. These last are also widely accepted
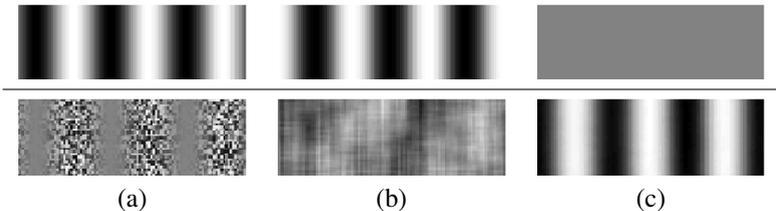
Figure 1: (a) First–Order (top) and Second–Order (bottom) sinusoidal modulations charac-terised by local variations in luminance and contrast respectively; (b) response of a standard (Intensity) Ranklet on the signals in (a); (c) response of our proposed Variance Ranklet (ver-tical orientation, filter size matched to wavelength of stimulus).

as a model of low-level human vision, namely the response of simple cells in the primary visual cortex [6]. Linear filters are in fact a natural choice to model sensitivity to intensity modulations (so called first–order stimuli), to which also Ranklets respond. The human visual system, however, also shows a similar response to contrast or variance modulations (second–order stimuli - Figure 1a). These are thought to be important, for instance, in the perception and discrimination of textures [9].

Contrary to first–order stimuli, the response to second–order stimuli is not easily mod-elled by linear filters. Perhaps the most common approach in the psycho-physical literature is the linear-nonlinear-linear (LNL) model, that consists of two linear filter stages separated by a non-linearity stage [11]. Features designed for second–order stimuli are rare in the ap-plied literature. Local variance estimators are indeed common, but these quantify variance itself rather than how (and especially along which orientation) it is modulated. Interestingly, even when non-parametric texture descriptors such as Local Binary Patterns are used, the variance measures employed are not themselves rank-based [20].

We present what we believe is the first set of rank features for the orientation selective detection of second–order stimuli. Variance Ranklets are defined in analogy with Ranklets, using the Siegel-Tukey statistics for dispersion instead of the Wilcoxon statistics for trans-lation. Both statistics are computable with a sorting operation and a sum and have the same null distribution, leading to a uniform treatment of information from first– and second–order stimuli. Variance Ranklets present the same pattern of orientation selectivity that charac-terises standard (Intensity) Ranklets and Haar wavelets.

We validate our features with texture classification experiments over a subset of the Vis-Tex [16] and of the Brodatz [5] databases. Starting from the responses of Intensity and Variance Ranklets, we compute the statistical descriptors described in [24] and classify the resulting feature vectors with a simple Nearest Neighbour (NN) classifier. We show that using Variance Ranklets together with Intensity Ranklets leads to significantly lower error rates than those obtained with standard Ranklets alone. We compare our recognition rates with recently published results obtained using four different algorithms [3, 13, 15, 19].

The rest of this paper is organised as follows: in Section 2 we give the relevant back-ground on Rank Statistics. After reviewing standard Ranklets in Section 3, we introduce Variance Ranklets in Section 4. An application to texture classification is discussed in Sec-tion 5 followed by experimental results in Section 6. Finally, conclusions are presented in Section 7.

2 Rank Statistics and Pixels

One way of describing local variations in images is to look at the statistical properties of neighbouring image regions. In this work, we use Rank Statistics to compare adjacent sets of pixels based on the intensity measurement performed at each image location. In accordance with the statistical literature, we denote the two sets of pixels as the Treatment and Control samples; their geometric arrangement on the image (that determines the orientation–selective behaviour of the filters) is specified in Section 3 below.

The point with Rank Statistics is to rely only on the relative order of the measurements rather than on their specific value. We therefore sort intensity values and replace them with their rank: for example, for 4 Treatment and 4 Control observations corresponding to intensity values $T = \{64, 128, 12, 56\}$ and $C = \{10, 75, 25, 100\}$ we have

Intensity	10	12	25	56	64	75	100	128
Sample	C	T	C	T	T	C	C	T
Rank	1	2	3	4	5	6	7	8

If we are now interested in knowing whether the observations in T are significantly larger or smaller than those in the C set, a natural choice is to use the Wilcoxon statistics W, defined as the sum of treatment ranks: in our case $W = 2+4+5+8$. Clearly a large value of W will indicate that the pixels represented in set T are on average brighter than those in set C, and conversely for a small value of W. The Wilcoxon statistics has been used in Ranklets (see Section 3 below) and implicitly, as shown in [22], in the rank transform [25]. Technically, the Wilcoxon statistics allows testing the null hypothesis that the observations in T and C are drawn from identical (though unspecified) distributions against the alternative that one of the distributions is translated.

If we are interested in testing against the alternative that the pixels in T differ from the pixels in C in *variability*, we can assign ranks in the following way:

Intensity	10	12	25	56	64	75	100	128
Sample	C	T	C	T	T	C	C	T
Rank	1	4	5	8	7	6	3	2

The sum of the treatment ranks, $ST = 4+8+7+2$, is known as the Siegel–Tukey statistics [21]. Clearly, high values of the statistics indicate that the observations in the T set are grouped more tightly around their median than the observations in the C set. Thus the Siegel–Tukey statistics can be used to test the data for spread, under the assumption that the two samples are drawn from distributions with equal median.

Other ranking schemes for testing against dispersion have been suggested — notably by Ansari and Bradley [2]. These tests have however been shown to be equivalent [10]. Also, all such tests require some assumption on the kind of admissible distribution pairs to be compared, for instance that they should have the same median [18]. Our choice of the Siegel–Tukey statistics is motivated by the fact that it has the same null distribution as the Wilcoxon statistics [12, 21], thus ensuring a more uniform treatment of the two kinds of stimuli that we consider.

We have until now made an implicit assumption that no two pixels in the local window of interest have the same intensity. When ties occur, these can be treated by replacing the ranks of the tied pixels with their average (or *midrank*). The statistical treatment of midranks is discussed in [12].
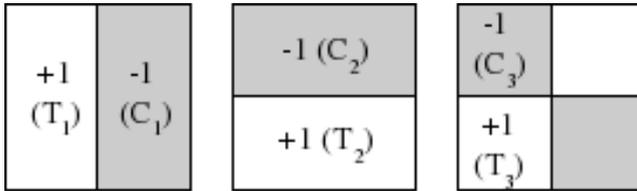
Figure 2: The three two-dimensional Haar wavelets $h_1(\vec{x})$, $h_2(\vec{x})$ and $h_3(\vec{x})$ (from left to right). Letters in parentheses refer to the "Treatment" and "Control" pixel sets (see Section 3).

3 Standard (Intensity) Ranklets

Ranklets (or, for the sake of clarity, Intensity Ranklets) are a family of multiscale rank filters based on the Wilcoxon statistics, that display an orientation selectivity pattern similar to Haar wavelets [22].

The orientation selective response is obtained by an appropriate choice of the T and C sets of observables (pixels intensity values) that are compared by the test. More specifically, in the case of Ranklets these sets are defined based on the three Haar wavelets $h_i(\vec{x}), i = 1, 2, 3$ supported on a given local window W, as sketched in Figure 2. The counter–images of $+1$ and -1 under these three functions provide three choices for T and C, namely $T_i = h_i^{-1}(\{+1\})$ and $C_i = h_j^{-1}(\{-1\})$.

Let $\pi(\vec{x})$ be a ranking of the pixels in W according to their intensity $I(\vec{x})$. The Wilcoxon statistics is computed separately for each of the three pairs of samples: $W^i = \sum_{\vec{x} \in T_i} \pi(\vec{x})$.

We centre W^i around its average value by computing $W_{YX}^i = W^i - (n+1)n/2$, where n is the number of pixels in T_i. The Mann-Whitney statistics W_{YX}^i is equal to the number of pixel pairs $(\vec{x}, \vec{y})$, with $\vec{x} \in T_i$ and $\vec{y} \in C_i$, such that $I(\vec{x}) > I(\vec{y})$ (for a proof, see [12]); its minimum and maximum values are thus 0 and n^2 respectively, with $n = \#T_i = \#C_i$. The standard Intensity Ranklets are defined as

$$R_W^i = 2W_{YX}^i/n^2 - 1, \tag{1}$$

so that their value ranges from -1 to $+1$ as the contrast between T_i and C_i is reversed. Thus Ranklets reproduce the orientation selectivity characteristics of the three Haar wavelets from which the T_i and C_i sets are derived.

The response of an Intensity Ranklet on a sample sinusoidal stimuli is shown in Figure 1(b). As can be seen, similarly to linear filters, Intensity Ranklets do not significantly respond to variance modulations.

4 Variance Ranklets

The definition of Ranklets given above can easily be generalised to second–order stimuli by replacing the Wilcoxon statistics W with the Siegel–Tukey statistics for spread introduced in Section 2.

This, however, requires a preprocessing step to ensure that the distributions of the two samples T_i and C_i have the same median. In practice, this is achieved by subtracting from the pixels in T_i and C_i the respective sample medians. The pixel values are then sorted

and the Siegel–Tukey ranks $\pi_{ST}(\vec{x})$ are computed in place of the standard ranks $\pi(\vec{x})$. The Siegel–Tukey statistics is then obtained as $ST^i = \sum_{\vec{x} \in T_i} \pi_{ST}(\vec{x})$, as illustrated in Section 2.

As ST^i and W^i have the same null distribution [21], we can centre the Siegel–Tukey statistics around its mean by subtracting the same constant used in Section 3 above, to give $ST_{YX}^i = ST^i - (n+1)n/2$; again, n is the number of pixels in T_i.

We then define the Variance Ranklets R_{ST} as

$$R_{ST}^i = 2ST_{YX}^i/n^2 - 1, \tag{2}$$

in analogy with Equation 1. The response of a Variance Ranklet also varies between -1 and $+1$, with the extremes indicating that the intensities of all the pixels in the T set are more tightly grouped around their median than those in the C set $(+1)$, or vice-versa for -1.

Figure 1(c) shows the response of a Variance Ranklet tuned to the vertical orientation on both first– and second–order stimuli. As can be seen, the filter responds selectively to variance modulations only.

5 Application: Texture Classification

We evaluate the effectiveness of the Variance Ranklets introduced above in an application to Texture Classification. This is motivated both by the importance of second–order stimuli for texture classification [9] and by the results recently obtained by Masotti and Campanini [15], that show that (Intensity) Ranklets compare favourably with Wavelets and Ridgelets in texture classification over a subset of the Vistex and Brodatz databases [5, 16]. For ease of comparison, we adopt the same choice of features and datasets as [15], although in combination with a simpler NN classifier as opposed to a Support Vector Machine (SVM). This is done to highlight the descriptive power of the features themselves rather than of the classifier.
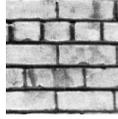
5.1 Feature Extraction

We evaluate the performance of our system using three sets of features, derived from Intensity Ranklets only, Variance Ranklets only and a combination of both. Feature extraction consists in rank-based filtering followed by quantisation and the extraction of statistical descriptors from the filter responses. These steps are described in detail below.

5.1.1 Filtering with Ranklets

We filter each texture image with a set of multiscale ranklets. Following [15] we use filters supported on a square window W with 7 different choices of the side of W, namely 4, 6, 8, 10, 12, 14, and 26 pixels. For each of these resolutions all 3 Ranklet orientations are computed (see Figure 2), yielding 21 filter responses for each of the Wilcoxon and Siegel–Tukey Ranklet types. We then take the absolute value of filter responses (to force insensitivity to contrast polarity) and quantise them into 11 equally spaced bins over the interval $[0,1]$.

Figure 3 shows the absolute value of the responses of Intensity and Variance Ranklets with various orientations and resolutions over a sample texture from the Brodatz album.

Texture image (D55, Brodatz database)

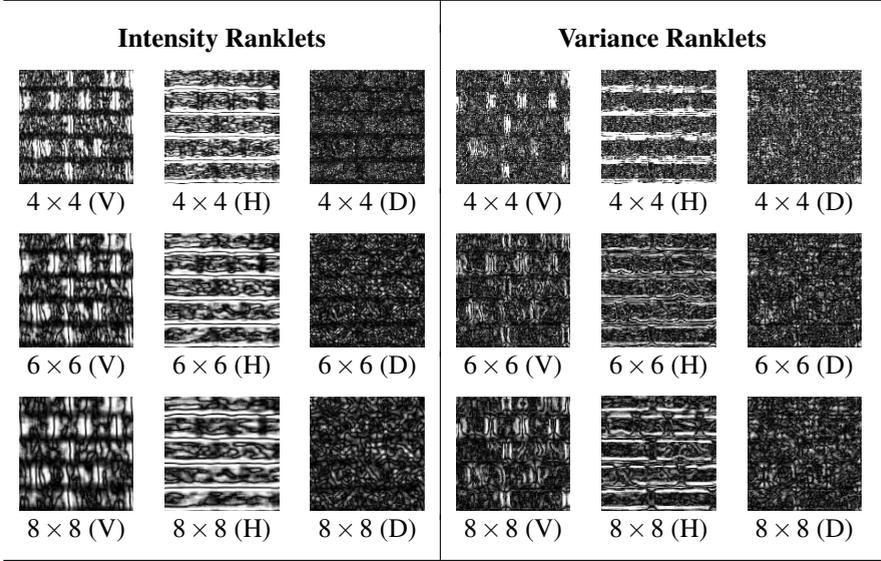| **Intensity Ranklets**			**Variance Ranklets**		
4×4 (V)	4×4 (H)	4×4 (D)	4×4 (V)	4×4 (H)	4×4 (D)
6×6 (V)	6×6 (H)	6×6 (D)	6×6 (V)	6×6 (H)	6×6 (D)
8×8 (V)	8×8 (H)	8×8 (D)	8×8 (V)	8×8 (H)	8×8 (D)

Figure 3: Absolute value of the response of Intensity and Variance Ranklets with the support size specified. Vertical, Horizontal and Diagonal orientations. Note the significant second–order component of the texture.

5.1.2 Statistical Descriptors

The quantised responses of each filter are used to derive a set of 11 statistical texture features that characterise the response of the filter on each image, as suggested by Xu et al. [24] (see also [15]). We define the histogram rh of Ranklet responses

$$rh(i) = \frac{n(i)}{\Sigma_{j=1}^{11} n(j)}, \qquad i = 1, ..., 11 \tag{3}$$

where $n(i)$ is the relative frequency of the i-th quantised Ranklet value. The co-occurrence matrix $rcm_{d,\theta}$, that encodes the probability of the transitions between every pair of coefficients as one moves across the image, is obtained as

$$rcm_{d,\theta}(i,j) = \frac{n_{d,\theta}(i,j)}{\Sigma_{l=1}^{11} \Sigma_{k=1}^{11} n_{d,\theta}(l,k)}, \qquad i, j = 1, ..., 11 \tag{4}$$

where $n_{d,\theta}(i,j)$ is the number of co-occurrences of Ranklet values quantised in bins i, j at a distance of d pixels along direction θ.

The first two feature descriptors are derived from the Ranklet histogram rh:

1. Mean Convergence

$$mc = \Sigma_{i=1}^{11} \frac{|rq(i).rh(i) - \mu|}{\sigma} \tag{5}$$

where μ and σ are the mean and standard deviation of the histogram rh and $rq(i)$ is the quantisation value corresponding to bin i;

2. Code Variance

$$cv = \Sigma_{i=1}^{11}(rq(i) - \mu)^2.rh(i) \tag{6}$$

The other 9 feature descriptors are derived from the mean of four Ranklet co-occurrence matrices $rcm_{d,\theta}$, where d is fixed to 1 and $\theta \in \{0°, 45°, 90°, 135°\}$, i.e: $rcm = \frac{1}{4}(rcm_{1,0°} + rcm_{1,45°} + rcm_{1,90°} + rcm_{1,135°})$. These are:

3. Code Entropy

$$ce = \Sigma_{i=1}^{11}\Sigma_{j=1}^{11}rcm(i,j).log_{10}(rcm(i,j)) \tag{7}$$

4. Uniformity

$$un = \Sigma_{i=1}^{11}\Sigma_{j=1}^{11}rcm(i,j)^2 \tag{8}$$

5-6. First– and Second–order element difference moments

$$fdm = \Sigma_{i=1}^{11}\Sigma_{j=1}^{11}|i-j|.rcm(i,j) \tag{9}$$

$$sdm = \Sigma_{i=1}^{11}\Sigma_{j=1}^{11}(i-j)^2.rcm(i,j) \tag{10}$$

7-8. First– and Second–order inverse element difference moments

$$fidm = \Sigma_{i=1}^{11}\Sigma_{j=1}^{11}\frac{1}{1+|i-j|}.rcm(i,j) \tag{11}$$

$$sidm = \Sigma_{i=1}^{11}\Sigma_{j=1}^{11}\frac{1}{1+(i-j)^2}.rcm(i,j) \tag{12}$$

9-11. Energy distribution of the ranklet co-occurrence matrix

$$edrcm1 = \Sigma_{i=1}^{3}\Sigma_{j=1}^{3}rcm(i,j) \tag{13}$$

$$edrcm2 = \Sigma_{i=1}^{6}\Sigma_{j=1}^{6}rcm(i,j) - edrcm1 \tag{14}$$

$$edrcm3 = \Sigma_{i=1}^{11}\Sigma_{j=1}^{11}rcm(i,j) - edrcm1 - edrcm2 \tag{15}$$

Applying the above descriptors to the quantised output of a filter on a texture image results in 11 features for each Ranklet filter. As proposed in [15], we force invariance to 90°-rotations by averaging the 11 descriptors across the 3 filter orientations displayed in Figure 2. Recalling that 7 scales are used for the filters, each feature image is described by a $7 \times 11 = 77$-dimensional feature vector derived from Wilcoxon Ranklets and a corresponding 77-dimensional vector derived from Siegel–Tukey Ranklets. These can be used either separately or jointly for classification, as described in Section 6.

Filters used	Classification Accuracy (%)			
	Vis-1	Vis-2	Bro-1	Bro-2
Standard Ranklets	96.04	91.33	96.15	99.73
Variance Ranklets	92.92	90.12	96.15	98.4
Standard + Variance Ranklets	**98.33**	**97.78**	**98.77**	**99.73**

Table 1: Classification accuracy over subsets of the VisTex and Brodatz databases, using a Nearest Neighbour classifier. Complementing the standard (Intensity) Ranklets with Variance Ranklets more than halves the error rate on three of the datasets.

5.2 Datasets

We present classification experiments over the four datasets used in [15]. Two of datasets, namely Vis-1[1] and Vis-2[2], comprised of 30 and 31 textures, are extracted from VisTex database [16]. Each texture image, of size 512×512 pixels, is decomposed into non-overlapping regions of size 128×128, resulting in 480 and 496 images respectively with every class (texture) represented by 16 samples.

The other two datasets, Bro-1[3] and Bro-2[4], comprised of 26 and 30 different textures respectively, are extracted from the Brodatz album [5]. These texture images, sized 640×640 pixels, are also decomposed into non-overlapping regions of size 128×128, yielding 650 and 750 images respectively; every class is here represented by 25 samples.

6 Experimental Results

We evaluate the descriptive power of the features by means of classification experiments using a simple NN classifier. The Euclidean distance is used for classification; tests are performed using leave-one-out (LOO) cross-validation across all 128×128 samples of all the different textures. Each dataset is considered separately.

Experimental results using feature vectors obtained from standard (Intensity) and Variance Ranklets either separately or jointly are displayed in Table 1 (information from the two types of ranklets is combined by concatenating the respective feature vectors into a 144-dimensional vector). As can be seen, features obtained from the Intensity Ranklets alone perform better than those obtained from the Variance Ranklets alone. This is in accordance with the findings of psycho-physical studies on the relative importance of the two types of stimuli the filters are designed to detect [9]. The best results are however obtained by combining the two sets of features: the addition of second order information more than halves

---

[1] Bark.0006, Brick.0000, Brick.0004, Brick.0005, Clouds.0001, Fabric.0000, Fabric.0006, Fabric.0007, Fabric.0013, Fabric.0015, Fabric.0017, Fabric.0019,Flowers.0005, Flowers.0006, Food.0000, Food.0001, Leaves.0003, Leaves.0012, Metal.0000, Metal.0002, Metal.0004, Misc.0001, Misc.0002, Sand.0000, Sand.0002,Stone.0005, Tile.0004, Tile.0008, Water.0005, Wood.0002

[2] Bark.0001, Bark.0004, Bark.0005, Bark.0006, Bark.0010, Brick.0000, Brick.0004, Brick.0005, Fabric.0000, Fabric.0002, Fabric.0007, Fabric.0009, Flowers.0005, Flowers.0006, Food.0001, Food.0002, Food.0004, Food.0005, Food.0006, Leaves.0002, Leaves.0003, Leaves.0004, Leaves.0008, Leaves.0010, Metal.0000, Metal.0005, Misc.0001, Sand.0002, Tile.0007, Water.0005, Wood.0002

[3] D1, D3, D4, D5, D6, D8, D9, D11, D16, D17, D18, D20, D21, D25, D26, D27, D28, D29, D33, D36, D78, D87, D92, D95, D101, D102

[4] D1, D3, D6, D11, D16, D17, D20, D21, D24, D28, D29, D32, D34, D35, D46, D47, D49, D51, D52, D53, D55, D56, D57, D65, D78, D82, D84, D85, D101, D104

the error rate on three of the datasets.

In Table 2 we report other recent published results referring to some of the datasets we use. These include approaches based on Ranklets and SVM [15], Wavelets and SVM [13], the Ridgelet transform and the L_1 distance [3] and invariant wavelet features in combination with NN [19]. Comparison with some of these works can only be indicative, as the size of the sub-images used for classification varies from 32×32 to 256×256 (128×128 is used here); however our results are quite comparable with the state-of-the-art and actually outperform some of the other algorithms on some of the datasets.

Published Algorithm	Classification Accuracy (%)			
	Vis-1	Vis-2	Bro-1	Bro-2
Standard Ranklets + SVM [15]	98.96	99.38	98.59	100.00
Ridgelets + L_1 distance [3]	99.00	-	-	-
Wavelets + NN [19]	-	99.52	93.55	-
Wavelets + SVM [13]	-	-	-	96.34
Our Approach	**98.33**	**97.78**	**98.77**	**99.73**

Table 2: Comparison with recently published results on the same datasets. See text for details of the differences between our approach and [15].

The most closely related approach is the work by Masotti and Campanini [15]. Apart from minor differences in the upper limit of the energy bands in Equations 14 and 15, they use the same statistical descriptors described in Section 5.1.2 over standard Ranklets in conjunction with SVM [15]. Their results are significantly better than those we obtain with standard Ranklets only (Table 1) due to their choice of the classifier (SVM vs our simpler NN). The significant reduction in error rate we obtain with Variance Ranklets suggests that by applying SVM classifiers to the combination of features extracted from Intensity and Variance Ranklets even better results would be achieved.

7 Conclusions

We introduced a new type of orientation-selective rank filters suitable for the detection of contrast modulations (second–order stimuli). Variance Ranklets are defined in analogy with the standard (Intensity) Ranklets, using the Siegel–Tukey statistics for dispersion in lieu of the Wilcoxon statistics for translation that appears in the original Ranklets. To the best of our knowledge, these are the first orientation selective rank–based features sensitive to variance modulations.

We validate our features with an application to texture classification, using texture images from two standard databases. Our experiments show that the use of Variance Ranklets together with Intensity Ranklets greatly improves classification accuracy with respect to standard Ranklets alone. Comparison with other recently published results shows that our features achieve low, state-of-the-art error rates when a simple Nearest Neighbour classifier is used.

In future work we will investigate the use of Variance Ranklets in combination with more sophisticated classifiers such as SVMs and their application to tasks for which Ranklets have already proved effective, such as pattern recognition on noisy images.

References

[1] T. Ahonen and M. Pietikäinen. A framework for analyzing texture descriptors. In *VISAPP (1)*, pages 507–512, 2008.

[2] A. R. Ansari and R. A. Bradley. Rank sum tests for dispersion. *Ann. Math. Statist.*, 31: 1174–1189, 1960.

[3] S. Arivazhagan, L. Ganesan, and T. G. Subash Kumar. Texture classification using ridgelet transform. *Pattern Recognition Letters*, 27(16):1875–1883, 2006.

[4] D. N. Bhat and S. K. Nayar. Ordinal measures for visual correspondence. In *Proceedings 1996 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 351–7, 1996.

[5] P. Brodatz. *Textures: a photographic album for artists and designers*. Dover, New York, 1966.

[6] J. G. Daugman. Two-dimensional spectral analysis of cortical receptive field profiles. *Vision Research*, 20:847–856, 1980.

[7] A. Del Bue, F. Smeraldi, and L. Agapito. Non-rigid structure from motion using ranklet-based tracking and non-linear optimization. *Image and Vision Computing*, 25 (3):297–310, 2007.

[8] M. Heikkilä, M. Pietikäinen, and C. Schmid. Description of interest regions with local binary patterns. *Pattern recognition*, 42:425–436, 2009.

[9] A. P. Johnson, N. Prins, F. A. A. Kingdom, and C. L. Baker Jr. Ecologically valid combinations of first– and second–order surface markers facilitate texture discrimination. *Vision Research*, 47:2281–2290, 2007.

[10] J. Klotz. Nonparametric tests for scale. *Ann. Math. Statist.*, 33:498–512, 1962.

[11] M. S. Landy and I. Oruc. Properties of second-order spatial frequency channels. *Vision Research*, 42(19):2311–2329, 2002.

[12] E. L. Lehmann. *Nonparametrics: Statistical methods based on ranks*. Holden-Day, 1975.

[13] S. Li, J. T. Kwok, H. Zhu, and Y. Wang. Texture classification using the support vector machines. *Pattern Recognition*, 36(12):2883–2893, 2003.

[14] M. Masotti. A ranklet-based image representation for mass classification in digital mammograms. *Medical Physics*, 33(10):3951–3961, 2006.

[15] M. Masotti and R. Campanini. Texture classification using invariant ranklet features. *Pattern Recognition Letters*, 29:1980–1986, 2008.

[16] MIT Media Lab. VisTex: vision texture image database. http://vismod.media.mit.edu/vismod/imagery/VisionTexture/, 1995.

[17] A. Mohammed, R. A. El-Khoribi, and L. Fekry. Discrete Hidden Markov Tree modelling of ranklet transform for mass classification in mammograms. *GVOP Special Issue on Mammograms*, 7:61–68, 2007.

[18] L. E. Moses. Rank tests of dispersion. *Ann. Math. Statist.*, 34(3):973–983, 1963.

[19] K. Muneeswaran, L. Ganesan, S. Arumugam, and K. Ruba Soundar. Texture classification with combined rotation and scale invariant wavelet features. *Pattern Recognition*, 38(10):1495–1506, 2005.

[20] T. Ojala, M. Pietikainen, and T. Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on PAMI*, 24(7):971–987, 2002.

[21] S. Siegel and Tukey. J. W. A nonparametric sum of ranks procedure for relative spread in unpaired samples. *J. Amer. Statist. Assoc.*, 55:429–445, 1960.

[22] F. Smeraldi. Ranklets: orientation selective non-parametric features applied to face detection. In *Proc. of the 16th ICPR, Quebec QC*, volume 3, pages 379–382, 2002.

[23] F. Smeraldi. A nonparametric approach to face detection using ranklets. In *Proceedings of the 4th International Conference on Audio and Video-based Biometric Person Authentication, Guildford, UK*, pages 351–359, 2003.

[24] R. Xu, X. Zhao, X. Li, and C. I. Chang. Target detection with improved image texture feature coding method and support vector machine. *International Journal of Intelligent Technology*, 1(1):47–56, 2006.

[25] R. Zabih and J. Woodfill. Non-parametric local transforms for computing visual correspondence. In *Proceedings of the 3rd ECCV*, pages 151–158, 1994.