TESTING A MULTIVARIATE MODEL FOR WAVELET COEFFICIENTS

Roland Kwitt, Peter Meerwald, Andreas Uhl

Department of Computer Sciences
University of Salzburg, Austria

Geert Verdoolaege

Department of Applied Physics
Ghent University, Belgium

ABSTRACT

In this paper, we introduce a Goodness-of-Fit test for the Multivariate Exponential Power (MEP) distribution, a multivariate extension of the Generalized Gaussian, which has recently gained considerable interest as a model for wavelet coefficients in the context of color image retrieval and spread-spectrum watermarking. We present a size and power study of this test and show Goodness-of-Fit results for wavelet coefficients of natural and texture images from various popular databases.

Index Terms— Multivariate modeling, Hypothesis testing, Wavelet transform

1. MOTIVATION

In the last few years, the MEP distribution has gained considerable interest in the community as a multivariate model for DWT coefficients of subbands on the same decomposition level and across color channels. In [1], Verdoolaege et al. exploited the MEP in the context of texture image retrieval, demonstrating a considerable retrieval performance increase compared to a marginal subband model, such as the Generalized Gaussian. In [2], we demonstrate similar performance improvements in the context of spread-spectrum watermarking where the MEP model proved particularly useful in capturing the association structure among DWT coefficients of different color channels. Formally, the MEP distribution is a particular case of the Kotz-type distribution, which in turn belongs to the family of elliptical distributions [3]. The probability density function (pdf) of the n-dimensional MEP distribution is given by

$$p(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{\Sigma},\beta) = \frac{n\Gamma\left(\frac{n}{2}\right)}{\pi^{n/2}\Gamma\left(1+\frac{n}{2\beta}\right)2^{1+\frac{n}{2\beta}}}|\boldsymbol{\Sigma}|^{-1/2} \exp\left\{-\frac{1}{2}\left[(\boldsymbol{x}-\boldsymbol{\mu})^T\boldsymbol{\Sigma}^{-1}(\boldsymbol{x}-\boldsymbol{\mu})\right]^{\beta}\right\}, \boldsymbol{x}\in\mathbb{R}^n \tag{1}$$

with shape parameter $\beta > 0$, location parameter (vector) $\boldsymbol{\mu} \in \mathbb{R}^n$ and positive definite (symmetric) $n \times n$ dispersion matrix $\boldsymbol{\Sigma}$. As β goes to infinity, the MEP model tends to a multivariate generalization of the uniform distribution, whereas for $\beta = 0.5$ we obtain a multivariate generalization of the Laplace distribution for instance. For $\beta = 1$, the MEP model reduces to the Multivariate Normal (MVN). Further properties can be found in [3]. In case of modeling DWT coefficients of natural (or texture) images, $\boldsymbol{\mu} = \boldsymbol{0}$ is a descent assumption and we adhere to the notation $\boldsymbol{X} \sim \mathrm{MEP}_n(\beta, \boldsymbol{\Sigma})$ to signify that a n-dimensional random vector $\boldsymbol{X}$ follows such a MEP distribution.

The MEP model is of particular interest in any application field where the assumption of a MVN distribution is either too rigorous or a too coarse simplification. Nevertheless, the reasoning for employing the MEP model is generally heuristic and, in the context of wavelet coefficient modeling, relies on the argument that the coefficients are highly non-Gaussian and exhibit leptokurtic histograms. Since the adaption of statistical tools to check the Goodness-of-Fit (GoF) of the MEP distribution is lagging behind so far, practitioners rely on visual inspection of the fit. This strategy, however, is only possible in one or two dimensions and can be quite misleading as well, especially due to the problem of choosing a reasonable bin size in case $n = 2$. To formally state the problem we tackle in this work, let $\mathcal{D} = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N\}$ be a collection of observations $\boldsymbol{x}_i \in \mathbb{R}^n$, e.g. wavelet coefficient vectors. Our objective is to test the hypothesis $\mathcal{H}_0 : \boldsymbol{X} \sim \mathrm{MEP}_n(\beta, \boldsymbol{\Sigma})$ against a general alternative ($\mathcal{H}_1$). The contribution of this work is a novel statistical tool to perform this hypothesis test. This eventually allows to quantify the assumption of a MVN or MEP model for wavelet coefficients. Methodologically, our approach is strongly influenced by the GoF test for multivariate normality of Smith and Jain [4] as well as an idea from the field of multivariate two-sample hypothesis testing (cf. [5]).

The remainder of this article is structured as follows: In Section 2, we discuss prior work on a three-stage GoF strategy and then introduce the main ingredients of our approach in Section 3. In Section 4, we present our experiments including a size and power study and some test results on real-world data; we conclude with Section 5.

2. ADAPTION OF PRIOR WORK

To the best of our knowledge, there exists no published GoF test for the MEP distribution, although Gomez et al. [3] sketch a possible test strategy in their seminal paper on the MEP distribution. Since, the MEP model is a particular type of elliptical distribution it admits a stochastic representation, i.e. if $\boldsymbol{X} \sim \mathrm{MEP}_n(\beta, \boldsymbol{\Sigma})$ then

$$\boldsymbol{X} \sim R\boldsymbol{A}^T\boldsymbol{U} \tag{2}$$

where the random vector $\boldsymbol{U} \in \mathbb{R}^n$ is uniformly distributed on the unit-sphere in the n-dimensional Euclidean space, $\boldsymbol{A}$ is an upper triangular matrix such that $\boldsymbol{\Sigma} = \boldsymbol{A}^T\boldsymbol{A}$, R is distributed according to the pdf

$$f_R(r;\beta) = \frac{n}{\Gamma\left(1+\frac{n}{2\beta}\right)2^{\frac{n}{2\beta}}} r^{n-1}\exp\left\{-\frac{1}{2}r^{2\beta}\right\}\boldsymbol{1}_{(0,\infty)}(r) \tag{3}$$

and R is independent of $\boldsymbol{U}$. Based on this representation and the moments of R (see [3]), the first stage of the GoF test is to check whether the $z_i = (\boldsymbol{x}_i^T\hat{\boldsymbol{\Sigma}}^{-1}\boldsymbol{x}_i)^{\hat{\beta}}$ ($\hat{\beta}$ denotes an estimate of β, see Section 3.1) follow a Gamma distribution with shape parameter 2 and scale parameter $n/2\hat{\beta}$, since $Z = (\boldsymbol{X}^T\boldsymbol{\Sigma}^{-1}\boldsymbol{X})^{\beta} \sim \Gamma(2, n/2\beta)$. We implement this check by means of a Chi-Square GoF test. In the second stage, we have to test whether the transformed sample

$$\boldsymbol{u}_i = \boldsymbol{y}_i/\|\boldsymbol{y}_i\| \text{ with } \boldsymbol{y}_i = \hat{\boldsymbol{\Sigma}}^{-\frac{1}{2}}\boldsymbol{x}_i, \quad i = 1,\ldots,N \tag{4}$$

This research project was funded by the Austrian Science Fund (FWF) under project no. L366-N15.

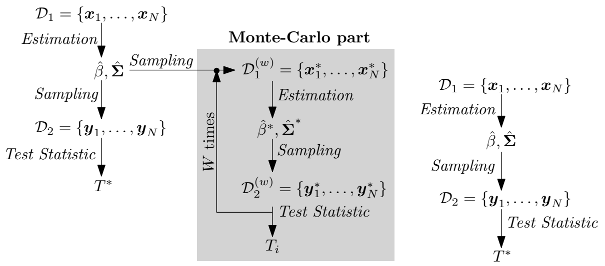

(a) *Monte-Carlo* variant (b) *Normal* variant

Fig. 1. Illustration of the two *variants* of our proposed GoF test.

is uniformly distributed on the unit sphere in $\mathbb{R}^n$ ($\hat{\boldsymbol{\Sigma}}$ denotes an estimate of $\boldsymbol{\Sigma}$). We perform this task by means of a Rayleigh test for uniformity on the sphere, originally proposed by Mardia and Rupp [6]. In the last step, we have to test if the random variable R is independent of the random vector $\boldsymbol{U}$. For that purpose, we employ a recently proposed test by Gretton et al. [7].

At that point, we notice that each stage strongly depends on the particular type of test that is used. However, the most crucial part is the fusion step where we have to combine the test outcomes of each stage into a final test result. Since no fusion strategy is given in [3], we choose the pragmatic (yet, rather strict) way to reject the overall null-hypothesis $\mathcal{H}_0$ in case just one stage shows evidence against its null-hypothesis. In Section 4.1, we conduct a size and power study of this test (which, to the best of our knowledge, has not been done so far) and compare it to our proposed approach.

3. INGREDIENTS OF OUR NEW APPROACH

Basically, we take the GoF approach of Smith and Jain [4] to test multivariate normality and adapt the relevant parts. The components of the test procedure are outlined in Fig. 1, where part (a) shows the Monte-Carlo variant of the test which is based on a Monte-Carlo estimate of the p-value and part (b) shows the second variant which relies on the asymptotic distribution of the test statistic, given $\mathcal{H}_0$ is true. In our setup, null-hypothesis is that the sample $\mathcal{D}_1$ is drawn from a MEP distribution with parameters β and $\boldsymbol{\Sigma}$. As we can see from Fig. 1, the critical parts of the GoF test are the *estimation* part, the *sampling* part and the selection of a suitable *test statistic*.

3.1. Parameter Estimation

Gomez et al. [3] mention moment matching as a suitable parameter estimation method. In [1], Verdoolaege et al. propose to use Maximum-Likelihood (ML) estimation. Moment-matching is described in [2]. Due to space limitations, we only focus on the ML case in the following.

In order to determine ML estimates for β and $\boldsymbol{\Sigma}$, we first formulate the log-likelihood expression $L(\beta, \boldsymbol{\Sigma}, \mathcal{S})$ of observing a random sample $\mathcal{S} = \{\boldsymbol{z}_1, \ldots, \boldsymbol{z}_N\}$ of realizations of N i.i.d. copies of a random variable $\boldsymbol{Z} \sim \mathrm{MEP}_n(\beta, \boldsymbol{\Sigma})$ as

$$
\begin{aligned}
L(\beta, \boldsymbol{\Sigma}; \mathcal{S}) &= N \log \Gamma\left(\frac{n}{2}\right) - N \log \Gamma\left(\frac{n}{2\beta}\right) + N \log(\beta) \\
&- \frac{Nn}{2}\log(\pi) - \frac{Nn}{\beta}\log(2) - \frac{N}{2}\log(|\boldsymbol{\Sigma}|) - \frac{1}{2}\sum_{i=1}^{N} \boldsymbol{v}_i^{\beta}
\end{aligned} \tag{5}
$$

with $\boldsymbol{v}_i \equiv \boldsymbol{z}_i^T \boldsymbol{\Sigma}^{-1} \boldsymbol{z}_i$. We then determine the partial derivatives of Eq. (5) w.r.t. β and $\boldsymbol{\Sigma}$

$$
\frac{\partial}{\partial \beta} L = N + \frac{Nn}{2\beta}\left[C + \psi\left(\frac{n}{2\beta}\right)\right] - \frac{\beta}{2}\sum_{i=1}^{N} \boldsymbol{v}_i^{\beta} \log(\boldsymbol{v}_i) \tag{6}
$$

and

$$
\frac{\partial}{\partial \boldsymbol{\Sigma}} L = -\frac{N}{2}\boldsymbol{\Sigma}^{-1} + \frac{\beta}{2}\sum_{i=1}^{N} \boldsymbol{v}_i^{\beta-1} \boldsymbol{\Sigma}^{-1} \boldsymbol{z}_i \boldsymbol{z}_i^T \boldsymbol{\Sigma}^{-1} \tag{7}
$$

with $C = \log(2)$ and ψ denoting the Digamma function. Setting the right-hand side of Eq. (7) to zero leads to

$$
\boldsymbol{\Sigma} = \frac{\beta}{N}\sum_{i=1}^{N} \boldsymbol{v}_i^{\beta-1} \boldsymbol{z}_i \boldsymbol{z}_i^T \tag{8}
$$

after some straightforward manipulations. Eq. (8) now allows to employ the method of successive approximation, starting with $\hat{\boldsymbol{\Sigma}}^{(1)} = \boldsymbol{I}$ until convergence is reached. Since we need to estimate β and $\boldsymbol{\Sigma}$ simultaneously, we solve ${}^{\partial}/_{\partial\beta} L = 0$ in each iteration of the successive approximation algorithm by using bisectioning in the interval $[0.1, 5]$[1], starting with $\hat{\beta}^{(1)} = 0.5$. Convergence is reached when $\sum_i \sum_j |\hat{\sigma}_{ij}^{(k)} - \hat{\sigma}_{ij}^{(k-1)}| + |\hat{\beta}^{(k)} - \hat{\beta}^{(k-1)}| < \epsilon$, where $\hat{\sigma}_{ij}^{(k)}$ and $\hat{\beta}^{(k)}$ denote the estimates in the k-th iteration.

3.2. Sampling

To draw a random sample from a MEP distribution with desired parameters β and $\boldsymbol{\Sigma}$, we can rely on Eq. (2). First, we draw a random sample $\boldsymbol{u}_1, \ldots, \boldsymbol{u}_N$ from the uniform distribution on the unit sphere in $\mathbb{R}^n$. We then perform a Cholesky decomposition of $\boldsymbol{\Sigma}$ to obtain $\boldsymbol{A}^T$ and generate another random sample $r_1, \ldots, r_N$ from the distribution given in Eq. (3). Eventually, we use $\boldsymbol{y}_i = r_i \boldsymbol{A}^T \boldsymbol{u}_i$ to generate a MEP random sample $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_N$. Let us review these steps in more detail: To obtain $\boldsymbol{u}_i$, we use the fact that the MVN distribution is radially symmetric. Drawing a random vector $\boldsymbol{u}_i$ from a standardized MVN distribution $\mathcal{N}(\boldsymbol{0}, \boldsymbol{1})$ and normalizing each vector element by $(\sum_j u_{ij}^2)^{1/2}$ consequently gives the desired result. The process of generating the random sample $r_1, \ldots, r_N$ is slightly more involved, though. In order to use the classic *inversion* method, we first determine the quantile function F_R^{-1} (i.e. the inverse cdf) corresponding to the pdf in Eq. (3) as

$$
F_R^{-1}(u; \beta) = 2^{\frac{1}{2\beta}} \left[P_u^{-1}\left({}^{n}/_{2\beta}, 1-u\right)\right]^{\frac{1}{2\beta}}, u \in [0,1] \tag{9}
$$

where $P_u(a, x)$ denotes the regularized (upper) incomplete Gamma function. We can then generate r_i using $r_i = F_R^{-1}(u_i; \beta)$ with u_i drawn from a uniform distribution on $[0, 1]$.

3.3. The Test Statistic

To find a suitable test statistic, Smith and Jain [4] borrow a test strategy from the field of two-sample hypothesis testing, where the objective is to test whether two samples $\mathcal{D}_1 = \{\boldsymbol{x}_1, \ldots, \boldsymbol{x}_N\}$ and $\mathcal{D}_2 = \{\boldsymbol{y}_1, \ldots, \boldsymbol{y}_N\}$ stem from the same underlying population without making any assumptions about the distribution family. The basic principle is to use some statistic of the pooled sample $\mathcal{D}_p = \{\mathcal{D}_1, \mathcal{D}_2\}, M := |\mathcal{D}_p| = 2N$ with known distribution under the null-hypothesis (of equal underlying population) to perform the test.

In our context, the sample $\mathcal{D}_1$ represents the collection of original observations, whereas the sample $\mathcal{D}_2$ is drawn from the assumed $\mathrm{MEP}_n(\beta, \boldsymbol{\Sigma})$ distribution with parameters fitted on the basis of $\mathcal{D}_1$.

[1] Assuming the true root β is in $[0.1, 5]$ is a reasonable assumption, at least for wavelet coefficients (cf. [1]).

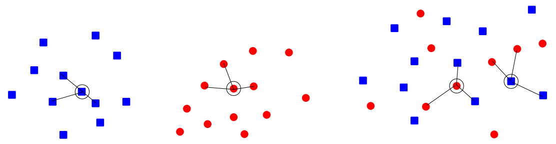

(a) $H \neq G,\ T = 6$ (b) $H = G,\ T = 2$

Fig. 2. Illustration of the two-sample hypothesis test proposed by Henze [8], based on the number of nearest neighbor coincidences for $k = 3$.

In [4], the assumed distribution is a MVN distribution and the authors rely on a two-sample test based on an Euclidean Minimum Spanning Tree. In this work, we rely on a similar strategy originally proposed by Henze [8]. The approach is based on the computation of the number of nearest neighbor (NN) coincidences, computed on $\mathcal{D}_p$ (using the Euclidean distance). A visualization of this idea is shown in Fig. 2, where H signifies the distribution of $\mathcal{D}_1$ (marked as blue squares) and G signifies the distribution $\mathcal{D}_2$ (marked as red discs). In both examples, we illustrate computation of the test statistic T, only considering two elements of each sample. The line of reasoning here is as follows: Given that the null-hypothesis (i.e. $\mathcal{H}_0 : H = G$) is true, we expect the number of NN coincidences to be low, cf. Fig. 2(b). On the other hand, if $\mathcal{H}_0$ is false, the number of NN coincidences will be high, cf. Fig. 2(a). We can apply the same reasoning in our context, only that $\mathcal{D}_2$ is now a random sample as defined above.

For a formal description of the test statistic, let $m : \mathcal{D}_p \to \{1, 2\}$ denote a function returning the sample membership of an element $\boldsymbol{z}_i$ of the pooled sample $\mathcal{D}_p$ and let $NN_i(r)$ denote the r-th nearest neighbor of $\boldsymbol{z}_i$. The test statistic $T_{k,M}$ for considering k nearest neighbors is

$$T_{k,M} = \frac{1}{Mk} \sum_{i=1}^{M} \sum_{r=1}^{k} \mathbf{1}_i(r) \tag{10}$$

where $\mathbf{1}_i(r)$ is 1 in case $m(\boldsymbol{z}_i) = m(NN_i(r))$ and 0 else. According to Schilling [5], we have the asymptotic (i.e. $M \to \infty$) result that in case $\mathcal{H}_0$ is true, the term $\phi = \sqrt{Mk} \cdot (T_{k,M} - \mu_{T_{k,M}|\mathcal{H}_0})/\sigma_{T_{k,M}|\mathcal{H}_0}$ follows a standard Gaussian distribution with $\mu_{T_{k,M}|\mathcal{H}_0} = \lambda_1^2 + \lambda_2^2$ and $\sigma^2_{T_{k,M}|\mathcal{H}_0} = \lambda_1\lambda_2 + 4\lambda_1^2\lambda_2^2(1 - C(2k,k)2^{-2k})$, where $C(a, b)$ denotes the binomial coefficient (i.e. a choose b) and $\lambda_i = N/M$ (i.e. in our case $\lambda_1 = \lambda_2 = 0.5$). Based on these expressions, it is straightforward to compute a p-value as $p = \mathbb{P}(T^* \geq T|\mathcal{H}_0) = 1 - F_T(\phi)$, where F_T denotes the cdf of the standard Gaussian distribution. Given a fixed significance level α, the null-hypothesis $\mathcal{H}_0$ is rejected in case $p < \alpha$.

Nevertheless, we have to bear in mind that the sampling procedure to generate $\mathcal{D}_2$ introduces some bias. This is obvious, since sampling is based on MEP parameters fitted on the basis of $\mathcal{D}_1$. The critical issue is that the NN coincidences test relies on the assumption of independent samples. Consequently, either size or power will be affected. To circumvent this problem, we follow the approach of [4] and estimate the critical region of the test using a Monte-Carlo approach, illustrated in Fig. 1(a). The iteration in the right branch of Fig. 1(a) is repeated W times and the p-value is estimated as $\hat{p} = (\#\{T_i \geq T^*\} + 0.5)/(W + 1)$. We denote the test variant based on Monte-Carlo p-value estimation as the *Monte-Carlo* variant and the test based on the asymptotic normality of $T_{k,M}$ as the *Normal* variant.

α	N	**Estimated $\hat{\alpha}$**		
		Gomez et al. [3]	*Monte-Carlo*	*Normal*
0.01	200	0.030	0.002	0.022
	400	0.028	0.001	0.002
	800	0.014	0.001	0.018
0.05	200	0.084	0.022	0.063
	400	0.118	0.012	0.014
	800	0.108	0.053	0.069
0.10	200	0.194	0.044	0.132
	400	0.212	0.026	0.048
	800	0.196	0.084	0.152

Table 1. Estimated size/level $\hat{\alpha}$ for the GoF test sketched by Gomez et al. in [3] as well as the two variants of our proposed GoF test, w.r.t. various levels of α and sample size N.

4. EXPERIMENTS

In order to assess the quality of the proposed GoF test(s) and our implementation of the test by Gomez et al., we conduct a study on the size, i.e. the test's probability of falsely rejecting $\mathcal{H}_0$, and power of the test. Regarding the experimental methodology, both size and power are evaluated by means of a Monte-Carlo simulation with $O = 500$ iterations for three-dimensional data (not to be confused with the *Monte-Carlo* variant of our test).

4.1. Size and Power Study

In each iteration of the Monte-Carlo simulation, we draw a random sample from $\text{MEP}_3(0.5, \boldsymbol{I})$ and determine the percentage of $\mathcal{H}_0$ rejections. Since the sample does actually stem from a MEP distribution, this percentage gives us an estimate $\hat{\alpha}$ for the test level/size. The size estimates for our implementation of the Gomez et al. test are obtained using the fusion strategy outlined in Section 2. Regarding the *Monte-Carlo* variant of our GoF test, we set the number of iterations W to 1000. Table 1 lists the estimated size/level $\hat{\alpha}$ for different sample size $N \in \{200, 400, 800\}$ and desired levels of $\alpha \in \{0.01, 0.05, 0.1\}$. For the GoF test of Gomez et al., we observe that $\hat{\alpha}$ is above the desired level α in all cases. Regarding the two variants of our proposed GoF approach, we see that the *Monte-Carlo* test is quite conservative, i.e. the percentage of false positives is always below the desired level. However, in case of the *Normal* test, the situation is different. Except for $N = 400$, the rejection rates are always slightly above the desired level.

To assess the power of the GoF tests, we sample from a two-component mixture of MEP distributions, i.e. $\pi_1 \cdot \text{MEP}_3(\beta_1, \boldsymbol{\Sigma}_1) + \pi_2 \cdot \text{MEP}_3(\beta_2, \boldsymbol{\Sigma}_2), \pi_1 + \pi_2 = 1$. We start from equal parameters $\beta_1 = \beta_2 = 0.5, \boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2$ and then gradually increase the shape parameter β_2 of the second mixture component. We refer to this scenario as *testing against shape alternatives*, presumably one of the most common scenarios in the context of wavelet coefficient modeling. The dispersion matrix $\boldsymbol{\Sigma}_1$ is chosen to resemble a real-world case and the mixture weights are $\pi_1 = \pi_2 = 0.5$. For each value of β_2 (and N), we perform O Monte-Carlo iterations and determine the number of $\mathcal{H}_0$ rejections.

Figure 3 shows the corresponding power plots. In case of the test by Gomez et al., we observe that our fusion strategy leads to reasonable power, even at moderate sample size. Regarding the two variants of our proposed test, both exhibit reasonable power as well with the *Normal* variant showing higher power again at moderate sample size. This can be explained by referring to Table 1, where the *Normal* variant is less conservative than the *Monte-Carlo* variant.

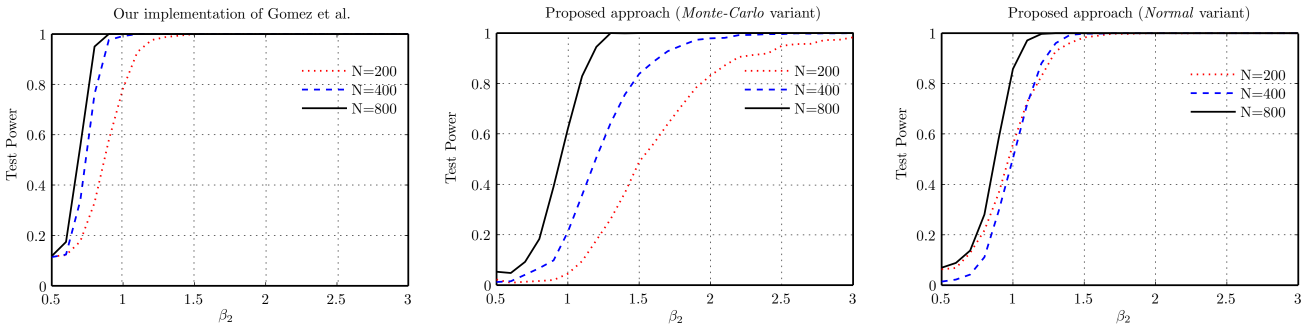

Fig. 3. Power plots ($\alpha = 0.05$) for the three GoF variants. The observations are sampled from a two-component MEP mixture where the shape parameter β_2 is gradually moved away from the shape parameter $\beta_1 = 0.5$ of the first mixture component.

Model	**Database**			
	Stex	Vistex (full)	Outex	UCID
$\text{MEP}_3(\beta, \mathbf{\Sigma})$	25.09	35.13	11.15	56.18
MVN (i.e. $\beta = 1$)	57.13	73.19	39.66	98.97

Table 2. $\mathcal{H}_0$ rejection rates for coefficients of (three level) DWT decomposed grayscale images at $\alpha = 0.05$.

4.2. Testing DWT Coefficients of Texture and Natural Images

We eventually turn to an actual application of the GoF test. We use our *Normal* test variant to check the GoF of a MEP_3 model for (horizontal,vertical,diagonal) DWT subband coefficients (using a three level decomposition with CDF 7/9 filters) of several database images. We use all images of the Stex[2], Vistex[3], Outex[4] and the UCID database[5]. To obtain equal power for each DWT decomposition level, we uniformly sample 500 coefficients from each subband and set $\alpha = 0.05$. In addition to full estimation of both MEP parameters, we perform a GoF test for the fix choice of $\beta = 1$, i.e. MVN, for comparative reasons. The $\mathcal{H}_0$ rejection rates are listed in Table 2. Compared to the MVN model, the MEP model is obviously the more suitable choice to capture wavelet coefficients statistics. The GoF results show a good fit for texture images and a slightly worse fit for natural images. Considering the simplicity of the MVN distribution (in terms of estimation complexity for instance), it is worth noting that in some cases the assumption of multivariate normality is actually not too coarse.

5. CONCLUDING REMARKS

We have introduced a novel statistical test to check the GoF of the MEP distribution, a popular multivariate model for DWT subband coefficients. We have further proposed a possible implementation of a prior GoF strategy and elaborated on MEP parameter estimation as well as sampling issues. MATLAB and C implementations of the tests will become available at `http://www.wavelab.at/sources` to foster reproducible research and make these tools available to the community.

Our size and power study demonstrates reasonable power of all proposed test strategies against shape alternatives and shows that the *Normal* variant exhibits controlled behavior in terms of test size. This is an attractive property, since the *Monte-Carlo* variant is far more computationally expensive. Our implementation of the GoF test by Gomez et al. exhibits remarkable power as well. However, we highlight that the employed fusion strategy is still somewhat arbitrary, whereas our proposed test(s) produce a *true* p-value and the notion of a significance level has a natural interpretation. We believe that efficient GoF tests for the MPE distribution are a valuable tool for many multivariate modeling tasks.

[2]Salzburg Textures, Online: `http://www.wavelab.at`
[3]MIT Vision Textures, Online: `http://vismod.media.mit.edu`
[4]Univ. of Oulo Textures, Online `http://www.outex.oulu.fi`
[5]Online: `http://vision.cs.aston.ac.uk/datasets/UCID`

6. REFERENCES

[1] G. Verdoolaege, Y. Rosseel, M. Lambrechts, and P. Scheunders, "Wavelet-based colour texture retrieval using the Kullback-Leibler divergence between bivariate Generalized Gaussian models," in *Proceedings of the IEEE International Conference on Image Processing (ICIP '09)*, Cairo, Egypt, Nov. 2009, pp. 265 – 268.

[2] R. Kwitt, P. Meerwald, and A. Uhl, "Color-image watermarking using multivariate power-exponential distribution," in *Proceedings of the IEEE International Conference on Image Processing (ICIP '09)*, Cairo, Egypt, Nov. 2009, pp. 4245–4248.

[3] E. Gomez, M-A. Gomez-Villegas, and J. M. Marin, "A multivariate generalization of the power exponential family of distributions," *Communications in Statistics – Theory and Methods*, vol. 27, no. 3, pp. 589–600, 1998.

[4] S. P. Smith and A. K. Jain, "A test to determine the multivariate normality of a data set," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, no. 5, pp. 757–761, Sept. 1988.

[5] M. F. Schilling, "Two-sample tests based on nearest neighbors," *Journal of the American Statistical Association*, vol. 81, no. 395, pp. 799–806, Sept. 1986.

[6] K. V. Mardia and P. Rupp, *Directional Statistics*, John Wiley and Sons Ltd., 2000, 2nd ed.

[7] A. Gretton, K. Fukumizu, C.H. Teo, L. Song, B. Schölkopf, and A.J. Smola, "A kernel statistical test of independence," in *Proceedings of the International Conference on Advances in Neural Information Processing Systems (NIPS'07)*, Vancouver, Canada, Dec. 2007, pp. 585–592.

[8] N. Henze, "A multivariate two-sample test based on the number of nearest neighbor type coincidences," *Annals of Statistics*, vol. 16, no. 2, pp. 772–783, June 1988.