

© IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

Efficient Texture Image Retrieval Using Copulas in a Bayesian Framework

Roland Kwitt, Peter Meerwald and Andreas Uhl

Abstract—In this article, we investigate a novel joint statistical model for subband coefficient magnitudes of the Dual-Tree Complex Wavelet transform which is then coupled to a Bayesian framework for Content-Based Image Retrieval. The joint model allows to capture the association among transform coefficients of the same decomposition scale and different color channels. It further facilitates to incorporate recent research work on modeling marginal coefficient distributions. We demonstrate the applicability of the novel model in the context of color texture retrieval on four texture image databases and compare retrieval performance to a collection of state-of-the-art approaches in the field. Our experiments further include a thorough computational analysis of the main building blocks, runtime measurements and an analysis of storage requirements. Eventually, we identify a model configuration with low storage requirements, competitive retrieval accuracy and a runtime behavior which enables the deployment even on large image databases.

Index Terms—Statistical Modeling, Copulas, Image Retrieval, Complex Wavelet Transform;

I. MOTIVATION

CONTENT-Based Image Retrieval (CBIR) applications have become increasingly popular in the last years. Two potential reasons for that are the vastly growing amount of digital image data and the free access to public image repositories. Both reasons are obviously connected to the dramatic price decline in the digital camera market. Systems which allow searching, sorting or retrieval of visual content have turned out to be essential in order to handle the vast amount of data. Since the holy-grail of image retrieval, i.e. searching by formulating semantic image queries, is still an open research problem [1], people resort to the paradigm of searching by example. In such a system, the user provides exemplary image content and gets a collection of *visually similar* images in return. It is further customary that there is an additional relevance feedback step at the end of the query process, where the user can mark relevant content and the system tries to incorporate the feedback to refine future query results. In this article, we base our work on the generic Bayesian CBIR formulation of Vasconcelos & Lippman, introduced in [2]. A schematic illustration of the main building blocks of this framework is shown in Fig. 1.

Since many images contain some sort of texture as part of the visual content, a substantial branch of CBIR research is

focused on texture characterization. Querying image content on the basis of texture similarity can be either part of a larger CBIR system, or a stand-alone system by itself. In the latter case, we think of online quality assurance in the manufacturing industry or real-time analysis of video frames in medical applications for example. A very popular concept in texture retrieval / classification research is to rely on a statistical characterization of filter outputs [3]. It has become common practice to extract features from transform coefficients of the Discrete Wavelet Transform (DWT) [4], [5], the Discrete Cosine Transform (DCT) [2], the Steerable Pyramid [6], [7], Gabor wavelets [8]–[10] or complex wavelet transforms [11]–[13]. A considerable amount of research work has been devoted to the statistical modeling of DCT or DWT subband coefficients, primarily focusing on the Generalized Gaussian (GG) distribution [5], [14], [15]. Recently, the Generalized Gamma (GGamma) distribution [16] and the Refined Histogram [17] have been proposed as reasonable alternatives. In particular, the Refined Histogram approach is based on the idea of modeling DWT coefficients by Product Benoulli distributions [18], [19] which can be estimated very efficiently. In addition to that, statistical inter- and intra-scale dependencies of DWT coefficients have been extensively studied [20], [21] and exploited for texture retrieval [22], estimation and segmentation [23], [24], predominantly by means of Hidden Markov Tree (HMT) models [25]. Recently, copula based modeling approaches to capture coefficient associations have gained considerable research interest [26]–[28] in the community. The contribution of this work resides in this last-mentioned research area.

To overcome the disadvantages of the real-valued DWT, such as missing phase information, lack of shift-invariance and lack of directional selectivity (see e.g. [29], [30]), complex wavelet transform variants have been proposed as reasonable alternatives. In this work, we consider one particular representative of this class, the Dual-Tree Complex Wavelet Transform (DTCWT) [29], which largely eliminates all aforementioned shortcomings at the cost of limited redundancy (i.e. four times in 2-D). Nevertheless, literature on the statistical characterization of complex transform coefficients is very limited. In summary, it can be stated that the Rayleigh, Weibull, or (Generalized) Gamma distribution are proposed to model the coefficient magnitudes [12], [31], [32], whereas inter-scale phase relationships are studied in [33] and a wrapped Cauchy or von Mises distribution is proposed to model the relative phase distributions [34]. In a recent work, we proposed a copula based modeling approach to capture the joint distribution of complex coefficient magnitudes across subbands of the same scale and different color channels and presented

Copyright (c) 2010 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

Roland Kwitt (corresponding author) and Andreas Uhl are with the Department of Computer Sciences, University of Salzburg, A-5020 Salzburg, Austria (e-mail: rk Witt@cosy.sbg.ac.at, uhl@cosy.sbg.ac.at, Tel.: +43 662 8044 6347, Fax: +43 662 8044 172). Peter Meerwald is currently with INRIA Rennes, France. The work was performed while he was with the University of Salzburg, Austria (e-mail: pmeerw@pmeerw.net).

Supported by the Austrian FWF Project No. 366-N15

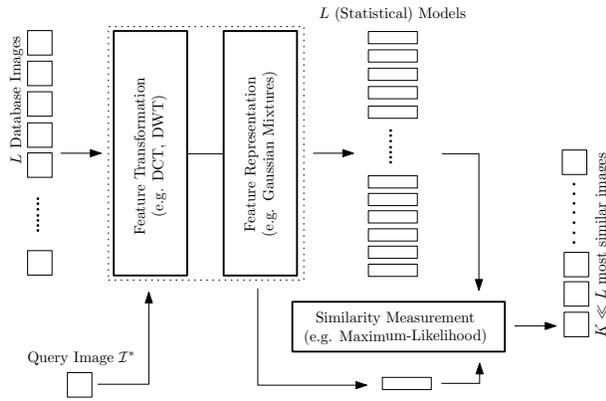


Fig. 1. Schematic illustration of the main building blocks of the Bayesian CBIR framework of Vasconcelos & Lippman [2].

promising preliminary texture retrieval results [35], mainly intended as a proof-of-concept study.

A crucial factor for the application of any CBIR system is computational efficiency. In a classic retrieval scenario – assuming that no hierarchical database model is in place – runtime is dominated by the query process. In particular, the time to process one image query is linear in the number of database images. Any parameters or features of the database images can be calculated at the time of storage. In the context of the Bayesian CBIR framework of Fig. 1, the computational bottleneck is the *similarity measurement* step. The two other building blocks, i.e. the *feature transformation* and the *feature representation* obviously contribute to the overall runtime, but in case of large databases the computational effort is negligible because we only need to compute the feature transformation once per query. Depending on the similarity measure, the feature representation might have to be computed once per query as well. We remark that a completely different challenge occurs when we consider the scenario of a limited set of database images and frequent image queries, i.e. matching video frames to a limited collection of templates. In this scenario, all three building blocks may present potential performance bottlenecks. In this work, however, we are only concerned with the aforementioned classic retrieval scenario.

A. Contribution

We investigate a novel, copula based statistical model for Dual-Tree Complex Wavelet Transform coefficient magnitudes which is (i) capable of capturing the association structure among transform coefficients of the same scale and different color channels and (ii) incorporates recent results [31], [36] on modeling marginal transform coefficient distributions. We substantiate the claims of previous texture retrieval research [28], [35] that Gaussian and Student t copulas are capable of capturing coefficient associations by presenting the first Goodness-of-Fit results in this context. In addition to that we establish a connection to the Bayesian, Minimum Probability of Error (MPE), image retrieval formulation of Vasconcelos & Lippman [2]. Our objective is to move beyond previous proof-of-concept studies and particularly address the trade-off

between retrieval rate and computational performance. In our point of view, the most prominent barrier for the widespread use of joint statistical models is the resulting complexity of the similarity measurement step. This study identifies a viable retrieval approach based on a Gaussian copula with Weibull margins.

Our experimental retrieval results on four texture image databases demonstrate (i) the superiority of copula based coefficient modeling compared to the exclusive use of marginal coefficient models (ii) that it is actually possible to achieve competitive runtime performance with respect to a collection of state-of-the-art approaches in the field. In consideration of a recent opinion article by Vandewalle et al. [37] on reproducible research, the full source code will become available at <http://www.wavelab.at/sources> to reproduce the experimental results.

B. Organization

The remainder of this article is structured as follows: in Section II, we briefly recapitulate the cornerstones of the MPE image retrieval formulation and then introduce the copula based modeling approach which includes a discussion on estimation and Goodness-of-Fit issues. In Section III, we investigate the computational requirements and present a simple strategy to enhance runtime performance. Section IV presents experimental retrieval results as well as a runtime study and an analysis of storage requirements. Finally, Section V concludes the paper with a discussion of the main points.

II. COPULA-DRIVEN IMAGE RETRIEVAL

A. Minimum Probability of Error Retrieval

In [2], [38], Vasconcelos & Lippman first introduced a Bayesian formulation of CBIR, also referred to as *Minimum Probability of Error (MPE)* retrieval. An image \mathcal{I} is considered to consist of a number of pixel observations $\mathbf{x} \in \mathcal{X}$ residing in the observation space \mathcal{X} . We assume that each image of the database belongs to one of M image classes, Y denotes a random variable with realizations in $\{1, \dots, M\}$ and p_Y is the corresponding probability mass function (p.m.f.). The first building block of the CBIR system is the *feature transformation* stage which is a mapping $t : \mathcal{X} \rightarrow \mathcal{Z}$ from the space of observations to the so called *feature space* \mathcal{Z} . The key issue here is to represent the image content in a domain which is more suitable for further processing. Accordingly, $\mathbf{z} = t(\mathbf{x})$ denotes a *feature vector*. The second building block is a probabilistic model describing how the feature vectors populate the feature space with respect to their class membership. The final part of the CBIR system is the retrieval function $g : \mathcal{Z} \rightarrow \{1, \dots, M\}$. In [2], the authors argue that the ulterior objective for designing this function is to minimize the probability of retrieval error. Given that the function $\omega : \mathcal{Z} \rightarrow \{1, \dots, M\}$ returns the true class membership of a feature vector \mathbf{z} , the objective is to minimize $\mathbb{P}(g(\mathbf{z}) \neq \omega(\mathbf{z}))$. From statistical classification theory (see e.g. [39]) we know that the function minimizing this criteria is the Bayes classifier

$$g(\mathbf{z}) = \arg \max_y p_Y(y|\mathbf{z}). \quad (1)$$

Following the formulation of [2], each image belongs to its own class with equal prior probability, i.e. $\forall y \in \{1, \dots, M\} : p_Y(y) = 1/M$. This is a reasonable simplification in CBIR, since it is hard to establish a-priori probabilities of database images. Hence, applying the Bayes rule to Eq. (1) leads to the Maximum-Likelihood (ML) selection criterion for image retrieval, i.e.

$$g(z) = \arg \max_y p_{Z|Y}(z|y) \quad (2)$$

which is substantially easier to handle than Eq. (1). We only have to compute the class-conditional likelihood $p_{Z|Y}(z|y)$ instead of the posterior probability $p_{Y|Z}(y|z)$. Consequently, the class-conditional probability density function (p.d.f.) $p_{Z|Y}$ constitutes the feature representation in our context. In any practical scenario $p_{Z|Y}$ has to be estimated from a collection of feature vectors z_1, \dots, z_N and the actual retrieval process will be based on a collection of query feature vectors z_1^*, \dots, z_K^* extracted from the query image \mathcal{I}^* . As we will later see, it is computationally beneficial to choose K smaller than N . To allow convenient and computationally tractable estimation of $p_{Z|Y}$, we assume that the feature vectors are i.i.d. and conditionally independent given their true class membership. In further consequence, this allows to write the ML selection rule as

$$g(z_1^*, \dots, z_K^*) = \arg \max_y \sum_{k=1}^K \log p_{Z|Y}(z_k^*|y). \quad (3)$$

Since each image belongs to its own class, we will omit Y in the notation $Z|Y$ from now on and instead indicate that a feature representation belongs to image \mathcal{I}_r by marking the model parameter (vector) $\Theta^{(r)}$. In the following sections, we are concerned with the development of a suitable feature representation $p_Z(z; \Theta)$.

B. Copula Based Statistical Subband Modeling

In the previous section, we introduced the three major building blocks of the MPE retrieval framework. For the choice of feature representation, we rely on previous research [31], [36] where we have successfully employed the DTCWT [40] in the context of image retrieval. Basically, a J -scale DTCWT leads to six complex detail subbands per decomposition level, capturing image details oriented along $\pm 15^\circ$, $\pm 45^\circ$ and $\pm 75^\circ$. Since we propose a novel joint model for DTCWT coefficient magnitudes, we are primarily concerned with the feature representation and the actual retrieval part of the MPE framework.

As the first issue of this section, we tackle the question whether we do actually need a joint statistical model for DTCWT coefficient magnitudes. In [31] for example, we relied on the assumption of total coefficient independency which enabled computationally efficient image retrieval. However, due to the fact that the DTCWT is a redundant transform (i.e. four-times redundant in 2-D), we have good reason to believe that a joint statistical model of coefficient magnitudes across scales and color channels will further enhance retrieval results. In fact, several authors have demonstrated that incorporating the association structure among transform coefficients considerably improved overall retrieval results [7],

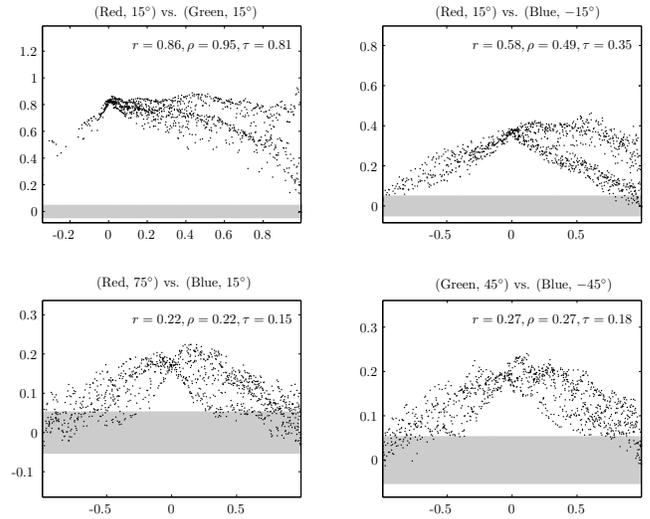


Fig. 2. Chi-Plots to illustrate the different types of association which can be observed between pairs of DTCWT coefficient magnitudes from different subbands and different color channels (on DTCWT decomposition level two).

[22], [41]. To get an impression about the pairwise association structure among coefficient magnitudes of different subbands we rely on a graphical tool commonly referred to as a Chi-Plot [42]. When considering two arbitrary subbands on one scale, the basic idea of a Chi-Plot is to transform pairs of coefficients in such a way that the resulting pairs (residing in the interval $[-1, 1] \times [-1, 1]$) reveal the structure of association. Consequently, a Chi-Plot can be considered as an extension of the classic scatter plot which is a customary tool to visualize possible dependencies. Nevertheless, scatter plots have the inherent disadvantage that it is hard to identify independent observations since the human visual system has problems to quantify the degree of randomness. In a Chi-Plot, however, departures from independency manifest themselves as deviations from the central region of the plot. A so called *tolerance band* is usually defined to allow slight scattering caused by sampling variability. Our implementation follows the algorithmic description given in [42]–[44], with the tolerance band enclosed by horizontal lines at $\pm c_p/\sqrt{N}$, where $c_p = 1.78$ and N denotes the number of observation tuples. In our Chi-Plots, the tolerance band is indicated as a gray-shaded region. Fig. 2 shows a set of exemplary Chi-Plots for pairs of DTCWT coefficient magnitudes from different subbands and color channels. The plots are chosen to visualize the most prominent types of association which can be observed for many DTCWT decomposed texture images. The strong deviation from the tolerance band, especially in the top two plots, further substantiates our conjecture that a statistical model which is capable of capturing the association structure may substantially improve the performance of any processing step which relies on a statistical characterization of the transform coefficients.

As the second topic of this section, we deal with the issue of incorporating existing information about the marginal distributions of the transform coefficient magnitudes. From previous research work [31], [32], [36] we know that the coefficient

magnitudes can be fairly-well modeled by either Weibull, Gamma or Generalized Gamma distributions. Obviously, we favor a joint statistical model which is capable of exploiting this information. We further do not want to constrain the joint model to only allow marginal distributions from the same distributional family (e.g. Gaussian). An elegant way to tackle both objectives at the same time is to exploit the mathematical construct of copulas. For a thorough theoretical introduction, we refer the reader to the classic textbooks by Joe [45] and Nelsen [46]. We only recapitulate the cornerstones of copula modeling next. Basically, a copula is a B -dimensional distribution function $C : [0, 1]^B \rightarrow [0, 1]$ with uniform marginals satisfying certain regularity conditions. In [47], Sklar showed that given a B -dimensional random vector $\mathbf{Z} = (Z_1, \dots, Z_B)$ with distribution function $F_{\mathbf{Z}}$ and marginal distribution functions F_{Z_1}, \dots, F_{Z_B} , there exists a B -dimensional copula C such that

$$F_{\mathbf{Z}}(\mathbf{z}) = C(F_{Z_1}(z_1), \dots, F_{Z_B}(z_B)), \mathbf{z} \in \mathbb{R}^B. \quad (4)$$

This theorem exploits the fact that every random variable can be transformed to a uniform random variable by applying the Probability Integral Transform (PIT) [48], i.e. the mapping $\mathbb{R}^B \rightarrow [0, 1]^B, \mathbf{z} = (z_1, \dots, z_B) \mapsto \mathbf{u} = (u_1, \dots, u_B)$ with $u_i = F_{Z_i}(z_i)$ and $\mathbf{u} \sim U([0, 1]^B)$. In other words, a copula can be considered as the distribution function of the PIT transformed margins. Since we will only consider random vectors \mathbf{Z} with continuous and strictly increasing marginal distribution functions, the copula C is uniquely determined on $[0, 1]^B$. As a corollary of Sklar's theorem [46], it follows that given a B -dimensional distribution function $F_{\mathbf{Z}}$ with margins F_{Z_1}, \dots, F_{Z_B} and copula C we have the relation

$$C(\mathbf{u}) = F_{\mathbf{Z}}(F_{Z_1}^{-1}(u_1), \dots, F_{Z_B}^{-1}(u_B)) \quad (5)$$

where the $F_{Z_i}^{-1}$ denote the quantile functions. This corollary substantially simplifies the process of finding a suitable statistical model for multivariate observations because the process of modeling the marginal distribution functions is completely decoupled from the process of modeling the association structure. As a consequence, this immediately allows to incorporate the knowledge we already have about the margins.

C. Copula Type & Estimation

We select two members of the family of elliptical copulas to capture the association structure among transform coefficients: the Gaussian copula and the Student t copula. Elliptical copulas arise from the family of elliptical distributions. In fact, they *are* the copulas of elliptical distributions and inherit all their properties such as simple generation of random numbers or well-known parameter estimation procedures for example. The copula of the multivariate Gaussian distribution with linear correlation matrix \mathbf{R} (i.e. $\text{diag } \mathbf{R} = \mathbf{1}$) is defined as

$$C(\mathbf{u}; \mathbf{R}) = \Phi(\Phi^{-1}(u_1), \dots, \Phi^{-1}(u_B)) \quad (6)$$

where Φ denotes the standard multivariate Gaussian distribution function and Φ^{-1} denotes the quantile function of

the standardized univariate Gaussian distribution. In the same manner, the Student t copula is defined as

$$C(\mathbf{u}; \mathbf{R}, \nu) = T_{\mathbf{R}, \nu}(t_{\nu}^{-1}(u_1), \dots, t_{\nu}^{-1}(u_B)) \quad (7)$$

where $T_{\mathbf{R}, \nu}$ denotes the standard multivariate Student t distribution, \mathbf{R} is defined as above, $\nu > 0$ denotes the degrees of freedom and t_{ν}^{-1} denotes the quantile function of the univariate Student t distribution. Since we will need the p.d.f.s $p_{\mathbf{Z}}$ of the copula based models in Section III, we recapitulate that the p.d.f. can be deduced from

$$p_{\mathbf{Z}}(\mathbf{z}) = \frac{\partial^B F_{\mathbf{Z}}(\mathbf{z})}{\partial z_1, \dots, \partial z_B} = \frac{\partial^B C(\mathbf{u})}{\partial z_1, \dots, \partial z_B} \quad (8)$$

which eventually leads to

$$p_{\mathbf{Z}}(\mathbf{z}) = \frac{\partial^B C(\mathbf{u})}{\partial u_1, \dots, \partial u_B} \prod_{i=1}^B \frac{\partial u_i}{\partial z_i} = c(\mathbf{u}) \prod_{i=1}^B p_{Z_i}(z_i) \quad (9)$$

since $\partial u_i / \partial z_i = \partial F_{Z_i}(z_i) / \partial z_i = p_{Z_i}(z_i)$. Hence, the p.d.f. of a copula based feature representation can be conveniently written as the product of the copula p.d.f. and the margin p.d.f.s. Based on the fact that the Gaussian copula and the Student t copula are the copulas of a multivariate Gaussian and multivariate Student t distribution, resp., it is straightforward to deduce the corresponding copula p.d.f.s. Considering Eq. (9), we just need to extract the term $c(\mathbf{u})$ from the p.d.f. $p_{\mathbf{Z}}$ of the multivariate Gaussian or multivariate Student t distribution. Given that \mathbf{Z} follows a standard B -variate Gaussian distribution with location vector zero and correlation matrix \mathbf{R} , we manipulate the p.d.f.

$$p_{\mathbf{Z}}(\mathbf{z}; \mathbf{R}) = \frac{1}{(2\pi)^{\frac{B}{2}} |\mathbf{R}|^{\frac{1}{2}}} \exp\left(-\frac{1}{2} \mathbf{z}^T \mathbf{R}^{-1} \mathbf{z}\right) \quad (10)$$

to get

$$c(\mathbf{u}; \mathbf{R}) = |\mathbf{R}|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \boldsymbol{\xi}^T (\mathbf{R}^{-1} - \mathbf{1}) \boldsymbol{\xi}\right) \quad (11)$$

with $\boldsymbol{\xi} = (\Phi^{-1}(u_1) \dots \Phi^{-1}(u_B))$. In a similar manner, we can determine the p.d.f. of the Student t copula. Given that \mathbf{Z} follows a B -variate Student t distribution with location vector zero, correlation matrix \mathbf{R} and $\nu > 0$ degrees of freedom, we manipulate the p.d.f.

$$p_{\mathbf{Z}}(\mathbf{z}; \mathbf{R}, \nu) = \frac{\Gamma\left(\frac{\nu+B}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right) (\nu B)^{\frac{B}{2}} |\mathbf{R}|^{\frac{1}{2}}} \left(1 + \frac{1}{\nu} \mathbf{z}^T \mathbf{R}^{-1} \mathbf{z}\right)^{-\frac{\nu+B}{2}} \quad (12)$$

to obtain

$$c(\mathbf{u}; \mathbf{R}, \nu) = \frac{|\mathbf{R}|^{-1/2} \Gamma\left(\frac{\nu+B}{2}\right) \left[\Gamma\left(\frac{\nu}{2}\right)\right]^B \left(1 + \frac{1}{\nu} \boldsymbol{\xi}^T \mathbf{R}^{-1} \boldsymbol{\xi}\right)^{-\frac{\nu+B}{2}}}{\left[\Gamma\left(\frac{\nu+1}{2}\right)\right]^B \Gamma\left(\frac{\nu}{2}\right) \prod_{i=1}^B \left(1 + \frac{\xi_i^2}{\nu}\right)^{-\frac{\nu+1}{2}}} \quad (13)$$

with $\boldsymbol{\xi} = (t_{\nu}^{-1}(u_1) \dots t_{\nu}^{-1}(u_B))$.

Regarding parameter estimation of the joint copula based model $p_{\mathbf{Z}}(\mathbf{z}; \boldsymbol{\Theta})$, we point out that it is computationally expensive and numerically cumbersome to simultaneously estimate the parameters of the marginal distributions and the copula parameters (denoted as the *exact* ML approach). Instead, we

follow a commonly-used two-step procedure termed the *Inference Functions from Margins (IFM)* method, introduced by Joe [45]. The IFM approach is based on a simple decoupling of the estimation procedure: given that we have an i.i.d. sample $\mathbf{z}_1, \dots, \mathbf{z}_N$, we first estimate the parameters $\theta_1, \dots, \theta_B$ of the margins (see [5], [49] for GG estimation and [36], [50] for Weibull / Gamma estimation details). Second, we use the obtained estimates $\hat{\theta}_j$ to perform the PIT on each \mathbf{z}_i and then estimate the copula parameters in a ML sense, i.e.

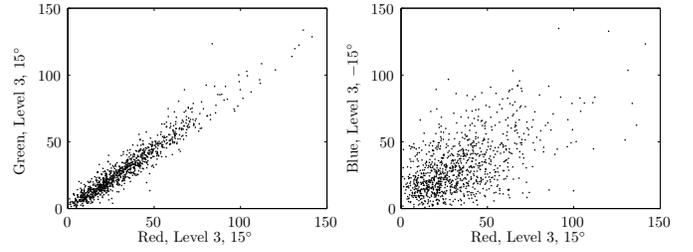
$$\hat{\Lambda} = \arg \max_{\Lambda} \sum_{i=1}^N \log c(F_{Z_1}(z_{i1}; \hat{\theta}_1), \dots, F_{Z_B}(z_{iB}; \hat{\theta}_B); \Lambda) \quad (14)$$

where $\Lambda = \{\mathbf{R}\}$ (Gaussian) or $\Lambda = \{\nu, \mathbf{R}\}$ (Student t), respectively. In case of the Gaussian copula, the ML estimate of \mathbf{R} is simply the sample correlation matrix of $\hat{\xi}_1, \dots, \hat{\xi}_N$ which can easily be shown by taking the first partial derivative w.r.t. \mathbf{R} of the log-likelihood function corresponding to Eq. (11) and setting the resulting term to zero. Unfortunately, the ML estimates of the Student t parameters \mathbf{R} and ν do not have such an explicit expression and must be determined by a numerical optimization algorithm. In this work we rely on MATLAB's `copulafit` routine to estimate ν and \mathbf{R} . Basically, the estimation algorithm employs numerical function minimization to find a minimum of the negative log-likelihood function corresponding to Eq. (13) w.r.t. ν . During this minimization procedure, \mathbf{R} is iteratively estimated using an algorithm proposed in a working paper by Bouyé et al. [51].

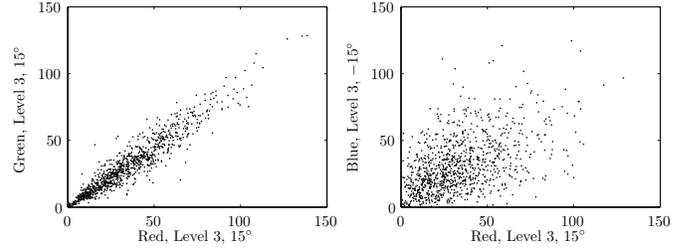
D. Quantifying the Goodness-of-Fit

To the best of our knowledge, there exists no commonly-accepted or recommended method to test the Goodness-of-Fit (GoF) of a copula model. Nevertheless, several approaches have been proposed recently in literature (see Genest et al. [52] or Berg [53] and references therein). The variety of ideas ranges from the reduction of the multivariate GoF problem to an univariate one (using the PIT), to parametric bootstrap procedures [54] or even the exploitation of positive definite bilinear forms [55]. In [35], we chose a pragmatic and straightforward approach, originally suggested by Genest and Favre [56], as a first step towards model selection. We plot pairs of original DTCWT transform coefficient magnitudes against random samples from fitted copula models. Two examples of such a plot are shown in Fig. 3 where we fit a Gaussian and a Student t copula with Weibull margins to two different subband combinations of the DTCWT decomposed *Bark.0008* texture image of the VisTex [57] database. The first row of Fig. 3 show a scatter plot of the original coefficient magnitudes, while rows two and three show 500 synthetically generated samples from the fitted copula models. The synthetic samples are generated by sampling from the copula and then transforming the uniform margins to Weibull margins by means of the Weibull quantile function.

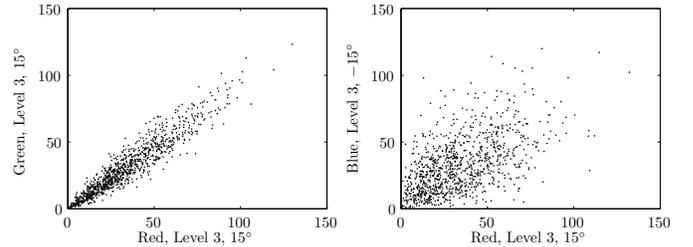
However, the large number of possible subband combinations limits the applicability of this approach to a preliminary visual inspections of model fit. To overcome this shortcoming,



(a) Original DTCWT transform coefficient magnitudes



(b) Student t (synthetic)



(c) Gaussian (synthetic)

Fig. 3. Scatter plots of (a) original DTCWT transform coefficient magnitudes compared against 500 samples drawn from fitted (b) Student t and (c) Gaussian copulas with Weibull margins.

we further experimented with the Akaike Information Criterion (AIC) [58] and the Bayes Information Criterion (BIC) [59] which both take into account the log-likelihood of the data under the given model and penalize additional parameters to avoid overfitting issues. Nevertheless, AIC and BIC are not adequate tools to address the problem of model selection in a hypothesis testing sense. They are rather useful as a means of selecting among possible candidate models without caring whether the models can actually describe the underlying data. To re-evaluate our selection of the Student t copula in [35], we implement a GoF test recently proposed by Genest et al. [52], [60]. The test is based on the computation of the Cramér-von-Mises statistic

$$\int_{[0,1]^B} C_B(\mathbf{u})^2 dC_B(\mathbf{u}) \quad \text{with} \quad C_B = \sqrt{N}(C_B - C_{\Lambda_B}) \quad (15)$$

where C_B denotes the empirical copula [46] and C_{Λ_B} denotes the estimated parametric copula under the null-hypothesis (i.e. either Gaussian or Student t). Regarding the actual implementation of the GoF test, we adhere to the parametric bootstrap algorithm [61] given in Appendix A of [52]. We choose 1000 bootstrap samples for our test. The null-hypothesis is rejected whenever the estimated p -value is lower than the significance level of 5%. Due to the fact that the parametric bootstrap

TABLE I
PERCENTAGE OF REJECTED NULL-HYPOTHESES OF A CHI-SQUARE GOF
TEST FOR THREE MARGINAL DTCWT DISTRIBUTION MODELS.

Database	Level	Margin Model		
		Weibull	Gamma	GGamma
VisTex (full)	1	11.68	8.87	11.12
	2	14.37	13.48	8.05
	3	14.81	14.91	9.45
ALOT	1	29.93	28.94	25.89
	2	25.20	26.52	13.13
	3	20.66	23.01	8.80
STex	1	19.90	12.93	24.01
	2	19.63	15.79	12.61
	3	18.70	17.68	10.06

procedure involves computation of C_{Λ_B} , cf. Eq. (15), we run into considerable computational problems since the test requires to compute multivariate Gaussian or multivariate Student t probabilities. This in turn requires computationally intensive multi-dimensional numerical integration for which we use the specifically-tailored algorithms presented by Genz [62] and Genz & Bretz [63]. As a consequence of the intensive computational demands, we limit our GoF study to 200 example textures of the VisTex (full) (see Section IV-A) database to get an impression of model fit. We select the subbands of DTCWT decomposition level three. Since we have three color channels and six subbands per scale, the joint statistical model is 18-dimensional. The rejection rates are 38.50% for the Student t copula and 35.50% for the Gaussian copula. Hence, we conjecture that both models are equally capable of capturing the association structure. In Section III, we will however see that the Gaussian copula is far more attractive from a computational point of view.

Nevertheless, the absolute rejection rates raise the question whether these rates are tolerable in the first place. Actually, it is hard to give a concrete answer to that question, since the rejection rate is strongly influenced by the power of the test and the significance level. Since we are dealing with a 18-dimensional dataset, the fact that test power increases with dimensionality leads to expect that the GoF test is quite conservative. Consequently, chances are high that the null-hypothesis gets rejected only because of minor deviations from the assumed model. In view of that, a rejection rate of 30–40% seems tolerable after all.

In order to evaluate the GoF of potential parametric margin models, i.e. Weibull, Gamma, Generalized Gamma, we perform Chi-Square GoF tests at the 5% significance level using all images of three texture databases (see Section IV-A). The bin width to compute the Chi-Square test statistic is chosen according to the standard setting used by the NIST software DATAPLOT [64], where the bin width is defined as $0.3s$, s denoting the sample standard deviation. In case of empty edge bins, the bins are combined with the next non-empty bin. Since the number of transform coefficients between two successive DTCWT decomposition levels differs by a factor of four, we uniformly sample 256 coefficients (without replacement) from each subband to achieve the same test power. This is reasonable, since we know that increasing the sample size likewise increases the power of a hypothesis

test. In case we omit this step, we would presumably blur the test results by risking higher rejection rates at the first decomposition levels. The exact test setup for one texture image is as follows: we decompose each (RGB) color channel separately by a three-scale DTCWT and then conduct the GoF test for the transform coefficient magnitudes of each of the 54 subbands. Table I lists the corresponding rejection rates per decomposition level. Based on these figures, we make the following observations: First, rejection rates are similar over the decomposition levels when considering the Weibull and Gamma distribution separately. Second, all three models seem equally capable of capturing the marginal coefficient statistics. We further observe that in case of the Generalized Gamma distribution, the rejection rates are less similar. We attribute this effect to numerical estimation issues related to the small (i.e. 256) number of coefficients that are used to estimate the three distribution parameters. Considering the fact that the Generalized Gamma distribution is computationally unhandy and does not significantly stand out with respect to the rejection rates, we restrict our choice of parametric margin model to the Weibull and Gamma distribution.

E. Wedding Copula Modeling and MPE Retrieval

For texture retrieval, we consider all available subbands of a specific decomposition level. In case of color images, a feature vector \mathbf{z} contains $B = 18$ components, where each component is a transform coefficient magnitude $z_i = |x_i|$, $x_i \in \mathbb{C}$ from one subband, i.e. $\mathbf{z} = (z_1, \dots, z_B)$. A concrete example of a copula based feature representation is a Gaussian copula with Weibull margins. Although it is reasonable to incorporate as much information as we can into the feature representation of each image, we run into problems when it comes to similarity measurement. In previous work [5], [16], [31], [65], it was shown that assuming independency among transform coefficients allowed to derive closed-form expressions for the Kullback-Leibler (KL) divergence between two feature representations. In case of copula based models however, no such closed-form expressions exist and we have to rely on alternative strategies. A first pragmatic approach we employed in [35] is to exploit the *Monte-Carlo* approximation of the KL divergence. Unfortunately, this approach has two inherent disadvantages: first, due to the dependency on random number generation, the KL divergence will differ to a certain extent (depending on the sample size) each time we compute the similarity between two feature representations. Second, the approach is computationally expensive since we need to estimate the joint statistical model for each query image, draw a random sample and compute the likelihood. A more reasonable way to perform the similarity measurement step is to employ the ML selection rule of the Bayesian CBIR framework, see Eq. (3). This is a natural choice, since it neither requires sampling nor parameter estimation of the query image's feature representation. Given a collection of query feature vectors $\mathbf{z}_1^*, \dots, \mathbf{z}_K^*$, the ML selection rule can be written as

$$g(\mathbf{z}_1^*, \dots, \mathbf{z}_K^*) = \arg \max_{r \in \{1, \dots, L\}} \sum_{i=1}^K \log p_{\mathbf{z}}(\mathbf{z}_i^*; \Theta^{(r)}) \quad (16)$$

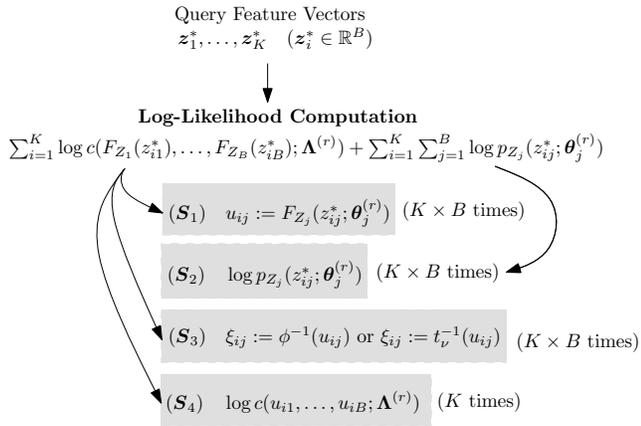


Fig. 4. Computational steps to compute the (log)-likelihood of a collection of query feature vectors under a copula based model from the database. The steps have to be repeated L times (i.e. for all models) to obtain the final retrieval result for query image \mathcal{I}^* .

where $\Theta^{(r)}$ denotes the parameter collection of the copula based feature representation of candidate image \mathcal{I}_r .

III. COMPUTATIONAL ANALYSIS

In MPE image retrieval, computation of the ML selection criterion is the key element to achieve competitive runtime performance. In particular, ML selection requires to evaluate the log-likelihood of a collection of query feature vectors under all L candidate models in the database. In the context of copula based feature representations, we can split the log-likelihood evaluation process into four parts, illustrated in Fig. 4 (assuming K query feature vectors). All four parts have to be carefully analyzed with respect to the computational requirements, since they present potential performance bottlenecks. The runtime measurements presented in this sections were performed using a C/C++ implementation of the ML selection rule on a 32-bit Intel Core2 Duo system clocked at 2.6 GHz with 4 GB of memory. The code was compiled using GCC 4.2 with options `-O2 -march=native -ffast-math`.

First, we discuss the issue of computing the logarithm of the Gamma / Weibull / GG p.d.f. and c.d.f., respectively, which covers steps (S₁) and (S₂) in Fig. 4. We include the GG distribution here, since we will present experimental retrieval results with a copula based feature representation of DWT coefficients using GG margins in Section IV as well. In case of the Weibull distribution with shape parameter $\alpha > 0$ and scale parameter $\beta > 0$, both the p.d.f. and c.d.f. have a closed-form expression. The logarithm of the p.d.f. is

$$\log p_Z(z; \alpha, \beta) = \log \alpha - \log \beta + (\alpha - 1) \log(z) - (\alpha - 1) \log(\beta) - z^\alpha \beta^{-\alpha} \quad (17)$$

whereas the corresponding c.d.f. has the form

$$F_Z(z; \alpha, \beta) = 1 - \exp \left\{ - \left(\frac{z}{\beta} \right)^\alpha \right\}. \quad (18)$$

We remark that only the terms $(\alpha - 1) \log(z)$ and z^α in Eq. (17) need to be calculated $K \times B$ times, all other terms can be pre-computed at the time of storage of the model, since they do not depend on the query image data. In case of the Gamma

distribution with shape parameter $\alpha > 0$ and scale parameter $\beta > 0$, the logarithm of the p.d.f. takes the form

$$\log p_Z(z; \alpha, \beta) = (\alpha - 1) \log(z) - \frac{z}{\beta} - \alpha \log(\beta) - \log \Gamma(\alpha) \quad (19)$$

and the c.d.f. is given by

$$F_Z(z; \alpha, \beta) = \frac{\gamma(\beta, z/\alpha)}{\Gamma(\alpha)} \quad (20)$$

where $\gamma(\cdot, \cdot)$ denotes the (lower) incomplete Gamma function [66]. Evidently, evaluation of the p.d.f. and c.d.f. is computationally more involved. Although most terms of Eq. (19) can be pre-computed, evaluation of the c.d.f. requires to compute the incomplete Gamma function. In our experiments, we use `gammq` implementation of [67] to compute Eq. (20). Last, the logarithm of the GG p.d.f. with shape parameter $\beta > 0$ and scale parameter $\alpha > 0$ is given as [68]

$$\log p_Z(z; \alpha, \beta) = \log(\beta) - \log(2\alpha) - \log \Gamma(1/\beta) - \left(\frac{z}{\alpha} \right)^\beta \quad (21)$$

and the corresponding c.d.f. has the inconvenient form

$$F_Z(z; \alpha, \beta) = \begin{cases} \frac{1}{2} \Gamma \left(\frac{1}{\beta}, \left(-\frac{z}{\alpha} \right)^\beta \right) \Gamma \left(\frac{\alpha}{\beta} \right)^{-1} & \text{if } z \leq 0 \\ 1 - \frac{1}{2} \Gamma \left(\frac{1}{\beta}, \left(\frac{z}{\alpha} \right)^\beta \right) \Gamma \left(\frac{\alpha}{\beta} \right)^{-1} & \text{if } z > 0 \end{cases} \quad (22)$$

where $\Gamma(\cdot, \cdot)$ denotes the (upper) incomplete Gamma function [66]. Similar to the Gamma distribution, evaluation of the c.d.f. is the numerically complex step. In particular, we exploit the $\Gamma(\cdot, \cdot)$ implementation `gammq` of [67] to compute the GG c.d.f. A direct runtime comparison of Eqs. (18), (20) and (22) reveals that calculation of the Weibull c.d.f. is approximately ten times faster than calculation of the GG and Gamma c.d.f., although this gap could possibly be closed by using a lookuptable approach to compute the Gamma functions.

Next, we turn to the computational effort required to evaluate of the Student t and Gaussian quantile functions, t_{ν}^{-1} and Φ^{-1} which covers step (S₃) in Fig. 4. In the Gaussian case, computation of Φ^{-1} can be accomplished using the inverse error function [66], i.e. $\Phi^{-1}(u) = \sqrt{2} \text{erf}^{-1}(2u)$ or by employing some fast approximation such as the Acklam algorithm¹ which uses two rational minimax approximations, one in the central region and one in the tails. We implement this approach for our experiments. In case of the Student t quantile functions, the situation is more complicated. For certain integer values of ν (e.g. $\nu = 1, 2$), t_{ν}^{-1} has an explicit expression. In general however, the fact that $\nu \in \mathbb{R}$ ($\nu > 0$) makes the evaluation of t_{ν}^{-1} less convenient. In our implementation, we use the `GSL`² routine `gsl_cdf_tdist_P` which computes the functional inverse of the Student t c.d.f. using the regularized Beta function [66]. Runtime measurements reveal a relative runtime difference of a factor ≈ 50 in favor of the Gaussian quantile function implementation.

¹See <http://home.online.no/~pjacklam/notes/invnorm/>.

²GNU Scientific Library, available at <http://www.gnu.org/software/gsl/>.

Finally, we address the effort to evaluate the logarithm of the Student t and Gaussian copula p.d.f.s, thus covering step (\mathcal{S}_4) in Fig. 4. When we consider the p.d.f. of the Gaussian copula, the only term in Eq. (11) we need to compute for each query feature vector is the quadratic form

$$-\frac{1}{2}\boldsymbol{\xi}^T(\mathbf{R}^{-1} - \mathbf{1})\boldsymbol{\xi}. \quad (23)$$

Obviously, \mathbf{R}^{-1} can be pre-computed at the time of model storage. Since we compute the inverse by means of a Cholesky decomposition, i.e. $\mathbf{R} = \mathbf{A}\mathbf{A}^T$, it is convenient to store the inverse of the lower triangular matrix \mathbf{A} instead of \mathbf{R}^{-1} . Given that $\mathbf{y} = \mathbf{A}^{-1}\boldsymbol{\xi}$, we can thus reformulate (23) as

$$-\frac{1}{2}\sum_{i=1}^B(y_i^2 - \xi_i^2). \quad (24)$$

In case of the Student t copula, the terms we need to compute for each query feature vector are

$$\left(1 + \frac{1}{\nu}\boldsymbol{\xi}^T\mathbf{R}^{-1}\boldsymbol{\xi}\right)^{\frac{\nu+B}{2}} = \left(1 + \frac{1}{\nu}\sum_{i=1}^B y_i^2\right)^{\frac{\nu+B}{2}} \quad (25)$$

and

$$-\left(\frac{\nu+B}{2}\right)\sum_{i=1}^B \log\left(1 + \frac{\xi_i^2}{\nu}\right) \quad (26)$$

given that \mathbf{y} is defined as above. Again, runtime measurements show that p.d.f. evaluation of the Gaussian copula is by a factor of ≈ 3 faster than p.d.f. evaluation of the Student t copula.

In consideration of the computational requirements to determine the log-likelihood function of the copula based feature representations, it is desirable to keep the number of query feature vectors K at a minimum. Similar to the query feature vector extraction strategy presented by Vasconcelos & Lippman in [2], we suggest a coefficient reduction step by random (uniform) sampling the subband coefficients, thereby reducing the data rate (and consequently the number of query feature vectors) by a factor of $1/n$. The speedup (cf. Fig. 6) will be proportional to the data reduction rate and considerably reduce the runtime of the ML selection process, however, reducing the number of coefficients might also negatively affect the retrieval rate (see Fig. 5). We will investigate this issue in the following section.

IV. EXPERIMENTS

Our experimental section covers two major issues: First, we evaluate whether the copula based subband models lead to competitive texture retrieval performance on four publicly-available databases in comparison to four state-of-the-art approaches in the field. Second, we address runtime as well as storage requirement issues and discuss the impact of the query feature vector reduction strategy.

A. Image Databases

We choose the Amsterdam Library of Textures (ALOT) [69], the MIT Vision Texture Database (VisTex) [57] and a

novel texture image database³ (Salzburg Textures, abbreviated as STex), captured by the authors under real-world conditions in the area around Salzburg/Austria for our experiments.

In case of the ALOT database, we only select the textures captured under the `c111` capture condition⁴. Regarding the VisTex images, we use the original 512×512 pixel versions of the textures available from the MIT Vision Texture website. There are 167 textures available from which we extract the same 40 image subset which appeared in various recent publications (see, e.g. [5], [16], [31], [70]). The full VisTex database will be denoted as VisTex (full), whereas the small subset will be denoted as VisTex (small). According to information on the website, VisTex images were captured under real-world conditions without studio lightning. Finally, the STex database consists of 476 images of different textures captured in the area around Salzburg / Austria using three cameras: a Canon IXUS 70, a Canon EOS 450D and a Nikon D40. Similar to the VisTex database, our image set is intended to resemble a real-life scenario.

B. Evaluation Criterion

To evaluate the performance of the retrieval system, we have to define a measure of retrieval *correctness*. We follow the common approach of counting the number of correct images among the top O retrieved images for each query, see [5], [71]–[73]. To state this measure in a more formal way, let $\mathcal{P}_1, \dots, \mathcal{P}_M$ denote the M parent images and let $\mathcal{I}_1, \dots, \mathcal{I}_L$ denote the texture images in the repository, obtained by splitting the original 512×512 pixel images into $S = 16$ non-overlapping 128×128 pixel blocks, i.e. $L = S \times M$. Further, we define a *parent indicator* function $p: \{1, \dots, L\}^2 \rightarrow \{0, 1\}$ as

$$p(i, j) := \begin{cases} 1, & \text{if } \mathcal{I}_i \text{ and } \mathcal{I}_j \text{ are splits of the same parent} \\ 0, & \text{else} \end{cases} \quad (27)$$

and let $R_j := \{r_1^{(j)}, \dots, r_L^{(j)} \mid r_1^{(j)} \leq r_2^{(j)} \leq \dots \leq r_L^{(j)}\}$ denote the index set of the sorted similarity values for the query image \mathcal{I}_j to all L candidate images (including the query itself). The percentage of correctly retrieved images for an arbitrary query image \mathcal{I}_j at the operating point of O retrieved images can be calculated as

$$s_O^{(j)} = \frac{1}{S} \sum_{i=1}^{O+1} p(j, r_i^{(j)}) \quad (28)$$

where the upper limit of the sum, $O + 1$, accounts for the fact that the query image is not excluded from the set R_j . This of course assumes that the query is naturally defined to be most similar to itself (which is always the case in our setup). The final retrieval rate t_O of the CBIR system at operating point O – calculated on the basis that each database image is used as a query once – can then be determined by

$$t_O = \frac{1}{SL} \sum_{j=1}^L \sum_{i=1}^{O+1} p(j, r_i^{(j)}). \quad (29)$$

³Available at <http://www.wavelab.at/sources/>.

⁴Depicted at <http://staff.science.uva.nl/~mark/ALOT/>.

Based on this evaluation setup, it is possible to construct Receiver Operating Characteristic (ROC) curves by plotting O against t_O . This allows to study the retrieval behavior as we increase the number of retrieved images (cf. Fig. 7). For practical purposes, reasonable values of O seem to be in the range of 16 to 40 images. As it is pointed out by Picard et al. [71], showing that a ROC curve of an approach lies above the ROC curve of another approach is a reasonable way to demonstrate a performance increase.

C. Comparison with State-of-the-Art Methods

To evaluate the competitiveness of the copula based subband models for texture image retrieval, we compare against a set of state-of-the art approaches which all basically follow the MPE paradigm of [2] (see Section II-A). The approaches can be grouped into two categories: those which are capable of handling color information innately and those which were originally proposed as grayscale-only approaches and thus require an artificial color extension. The first two approaches we briefly describe next are representatives of the first category, while the remaining three approaches are representatives of the second category.

First, we choose the original Embedded Multiresolution Mixture (EMM) model approach of Vasconcelos & Lippman [2] with eight mixture components and diagonal covariance matrices. Mixture model parameters are estimated by the classic EM algorithm [74] initialized by a modification (see [1]) of Gray's codeword splitting procedure [75]. The images are first converted to the YBR color model and a 8×8 window is shifted over the image by two pixel increments (in both directions) to obtain 48 DCT coefficients / window as feature vectors for parameter estimation (i.e. the first 16 DCT coefficients of each color channel are interleaved according to the pattern YBRYBR. . .). To extract query feature vectors, a non-overlapping 8×8 block DCT is used instead which finally gives 256 query feature vectors.

As a second approach, we implement a multivariate DWT coefficient modeling approach, presented by Verdoolaege et al. [70]. The authors propose a joint model for DWT coefficients based on the multivariate Power Exponential (abbreviated as multiv. Power. Exp.) distribution to capture dependencies of coefficients across color channels. The geodesic distance on the statistical manifold of multivariate Power Exponential distributions is then used as a similarity measure for retrieval. The distribution parameters are estimated by means of moment matching, as originally suggested by Gomez et al. [76].

Third, we implement the work of Do & Vetterli presented in [5]. The authors propose to model DWT coefficients by means of GG distributions and use a closed-form expression for the KL divergence during the retrieval process. The distribution parameters are estimated by means of the Newton-Raphson procedure proposed in the appendix of [5] (initialized by moment estimates, see [15]). To handle color images, we implement the straightforward extension to treat each color channel separately and finally sum up the KL divergences. This is also the extension we use for the following two approaches.

Fourthly, we test against an approach we presented in [31], [36]. The basic idea is similar to the work of Do & Vetterli [5], except that we employ the DTCWT for feature transformation and model the complex transform coefficient magnitudes by Weibull distributions to obtain a suitable feature representation. Parameters are estimated by means of ML estimation (initialized by moment estimates) [50] and the KL divergence minimization criterion is used during the actual retrieval process. As a matter of fact, this approach can be considered as a *lightweight* version of the proposal we present in this work, since coefficient associations are not incorporated.

As a final approach, we implement another extension to the work of Do & Vetterli [5], presented by Choy et al. [16]. The authors argue that the (three-parameter) Generalized Gamma distribution (GGamma) should be used to model DWT coefficient magnitudes and derive a closed-form expression to the KL divergence which is exploited to measure similarities during the retrieval process. Parameter estimation is accomplished in a ML sense by means of a globally-convergent algorithm proposed by Song [77].

For all approaches based on the DWT or DTCWT, the number of decomposition levels is set to three (as in [5], [70] for instance). Kingsbury's Q-Shift (14, 14)-tap filters are used for DTCWT decomposition levels greater than two in combination with (13, 19)-tap near-orthogonal filters for level one [78]. In case of the DWT we use CDF 7/9 filters.

Regarding the configuration of the copula based subband models, we test all combinations of Gaussian and Student t copulas with Weibull and Gamma margins. Further, we present experimental results for a Gaussian / Student t copula using GG margins as a multivariate extension of the work of Do & Vetterli [5]. This strategy has already been proposed by Stitou et al. [28], however, with the similarity measurement part based on a normalized l_1 norm between feature vectors composed of the copula and marginal distribution parameters. We prefer the ML selection criterion for retrieval, since ML selection has a thorough theoretical foundation, while the normalized l_1 distance is a rather arbitrary choice.

D. Retrieval Rate

In Table II, we summarize the retrieval rates at the operating point of $O = 16$ retrieved images for all four texture image databases. In addition, we list whether the retrieval results of the top approach per database (marked bold) show statistically significant differences to the results of the other approaches of the same database. For that purpose, we employ a paired sign-test at the 5% significance level. We choose this test, since it is a non-parametric test and it is further considered to be the only valid test to assess statistical significance in information retrieval experiments [79]. In case of a paired t-test for instance, we could not guarantee that the Normality assumption for the underlying population is met. In our setup, the test design is as follows: First, we determine the average retrieval rate per parent image \mathcal{P}_j (at $O = 16$ obviously), since we want to test on a per parent texture basis. These rates are computed by calculating the arithmetic mean over the

retrieval rates of the $S = 16$ query images corresponding to \mathcal{P}_j . Next, given M such rates for approach A and B , denoted by p_1^A, \dots, p_M^A and p_1^B, \dots, p_M^B , resp., the test statistic counts the number of times the difference $p_i^A - p_i^B, i = 1, \dots, M$ is positive or negative. In case the null-hypothesis of *no* statistically significant difference between A and B is true, the test statistic follows a binomial distribution. This eventually allows to determine a p -value. A rejection of the null-hypothesis is indicated by a superscript '**' in Table II, next to the retrieval rate t_{16} for each approach.

Regarding the effectiveness of joint statistical modeling in general, we note that the copula based approaches as well as the multiv. Power Exp. / Geodesic approach [70] show consistently higher retrieval rates than their relatives based on marginal subband models (e.g. used in [5], [16], [31]). We further observe that approaches based on the Student t copula exhibit very similar rates to approaches based on the Gaussian copula. This observation reflects the GoF test results of Section II-D and substantiates the claim that both copulas are equally capable of capturing coefficient associations. Regarding the choice of marginal distribution for the copula based feature representations, the situation is similar. Neither the Gamma nor the Weibull model particularly stand out with respect to the retrieval rate which again reflects the GoF test results of Table I. Nevertheless, from a computational point of view, the results are appealing since the least computationally expensive copula based feature representation (i.e. the Gaussian copula with Weibull margins) proves to be a good choice in practice. It even exhibits the highest retrieval rate on ALOT and is ranked among the top approaches on the remaining three databases. In fact, there is no incentive to use the Student t copula and the Gamma distribution here. Regarding the performance of the remaining approaches, we note that two competitors stick out particularly: First, it evident that the original Bayesian CBIR approach of Vasconcelos & Lippman (i.e. EMM / ML) – which is not specifically tailored for texture retrieval – is ranked among the top two results on the three large databases and ranked fourth on VisTex (small). Second, the multiv. Power Exp. / Geodesic approach of Verdoolaege et al. [70] exhibits similar performance, being ranked among the top two approaches on VisTex (small), ALOT, STex and ranked fourth on VisTex (full). Henceforth, we refer to these approaches as the two major competitors to our proposal.

As an interesting observation, we highlight that the analytically complex Generalized Gamma feature representation of Choy et al. [16] consistently exhibits slightly lower retrieval rates than the GG based feature representation of Do & Vetterli [5]. In some sense, the situation is similar to the discussion regarding the choice of copula, where we could not provide a convincing argument in favor of the more complex model. Even the objection that ML estimation is more complex in case of the GG distribution – as it is argued by Choy et al. – is not convincing, since Song [80] proposed a fast, globally convergent and computationally inexpensive Newton-Raphson iteration which is completely free of any kind of Gamma function.

In order to have a comparison to two non-parametric ap-

TABLE II
RANKING AND RETRIEVAL RATE ($O = 16$) ON EACH DATABASE.

Database	Approach	t_{16} [%]
VisTex (small) 40 Images	Multivar. Power Exp. / Geodesic [70]	91.2
	Gaussian Copula Weibull / ML	89.5
	Gaussian Copula Gamma / ML	89.1
	EMM / ML [2]	88.9
	Student t Copula GG / ML	88.9
	Student t Copula Gamma / ML	88.9
	Gaussian Copula GG / ML	87.5*
	Student t Copula Weibull / ML	86.9*
	DTCWT Weibull / KL [31]	84.0*
	DWT GG / KL [5]	82.0*
DWT GGamma / KL [16]	81.0*	
VisTex (full) 165 Images	Multivar. Power Exp. / Geodesic [70]	69.3
	EMM / ML [2]	67.7
	Student t Copula Gamma / ML	63.8*
	Student t Copula GG / ML	63.2*
	Gaussian Copula Weibull / ML	63.0*
	Student t Copula Weibull / ML	62.4*
	Gaussian Copula Gamma / ML	61.9*
	Gaussian Copula GG / ML	61.0*
	DTCWT Weibull / KL [31]	55.4*
	DWT GG / KL [5]	53.2*
DWT GGamma / KL [16]	53.1*	
ALOT 250 Images	Gaussian Copula Weibull / ML	54.1
	EMM / ML [2]	53.0
	Student t Copula GG / ML	50.8*
	Multivar. Power Exp. / Geodesic [70]	49.3*
	Gaussian Copula GG / ML	49.0*
	Student t Copula Gamma / ML	47.5*
	Gaussian Copula Gamma / ML	46.9*
	Student t Copula Weibull / ML	46.7*
	DWT GG / KL [5]	42.3*
	DWT GGamma / KL [16]	40.7*
DTCWT Weibull / KL [31]	40.6*	
STex 476 Images	EMM / ML [2]	73.7
	Multivar. Power Exp. / Geodesic [70]	71.3*
	Gaussian Copula Weibull / ML	70.6*
	Gaussian Copula Gamma / ML	69.4*
	Student t Copula GG / ML	65.6*
	Gaussian Copula GG / ML	65.2*
	Student t Copula Gamma / ML	64.3*
	Student t Copula Weibull / ML	63.1*
	DTCWT Weibull / KL [31]	58.8*
	DWT GGamma / KL [16]	52.9*
DWT GG / KL [5]	49.3*	

proaches as well, we implement the popular Local Binary Pattern (LBP) approach of [81] and the Multiresolution Histogram (MRH) approach proposed in [82]. In case of the LBPs, we use a standard 3×3 pixel neighborhood and employ the Histogram Intersection [83] metric to compute the similarity between two LBP feature representations. For the MRH approach, we use four pyramid levels of the Burt-Adelson pyramid, a Gaussian kernel for histogram smoothing and a fixed bin-width for the intensity resolution of the histograms (fixed to eight bit). As suggested in [82], the l_1 norm is used to measure the similarity between two MRH feature representations. For both approaches, the extension to cope with color images (i.e. RGB in our case) is straightforward by means of histogram concatenation. At the operating point of $O = 16$ retrieved images, the LBP approach exhibits 86.14% on VisTex (small), 59.19% on VisTex (full), 42.60% on ALOT and 58.77% on STex. The MRH approach exhibits 78.82% on VisTex (small), 49.61% on VisTex (full), 42.14% on ALOT and 49.94% on STex. We conclude that without any additional

effort to tune the performance of both approaches, they show up in the bottom half of the (per database) ranking listed in Table II. Nevertheless, we point out that computing the Histogram Intersection or the l_1 norm is substantially less expensive in terms of arithmetic operations than computing any of the similarity measures we use for the probabilistic feature representations.

E. Runtime Measurements & Storage Requirements

In this section, we take a closer look at the absolute retrieval runtime and storage requirements for the approaches of Section IV-C. All similarity measures, i.e. the ML selection criterion, the KL divergences and the geodesic distance, are implemented in C++. Numerical functions are provided by the GSL library and Numerical Recipes [67] as indicated in Section III. For matrix computation, we employ the Eigen template library for linear algebra⁵. We tried to achieve approximately the same level of optimization for all implementations: First, terms which occur repeatedly in an expression are calculated only once. Second, any matrix inversions or determinant calculations are performed offline, in case there is no dependency on the query image data. Third, all operations such as \log or $\log \Gamma$ that are independent of the query image data are computed once per query, and not for each query feature vector. We remark that these computations could be performed offline as well and then stored with the feature representation. Nevertheless, we omit this additional optimization step in order not to distort the pure storage requirements. We measure the average runtime of one image query to 1024 candidate images on a 32-bit Intel Core2 Duo system at 2.6 GHz with 4 GB of memory. Table III lists runtime numbers in milliseconds as well as the storage requirements for one feature representation of each approach (in bytes) using 64-bit double precision values (IEEE 754-1985). Note that the indicated runtime does *not* include computation of the feature transformation (and feature representation in case of [5], [16], [31], [70]) for the query image. Regarding the calculation of the exact storage requirements, we note that in case of symmetric $N \times N$ matrices, we only need to store $(N \times N + N)/2$ matrix components. Further, in case of [1], the covariance matrices of the Gaussian mixture model have diagonal form and in case of [70], the scatter matrix is a 27×27 block-diagonal matrix with symmetric block matrices of size 3×3 .

The KL similarity functions exhibit by far the fastest runtime in Table III, yet the retrieval rate results are ranked at the bottom for each database (cf. Table II). Computation of the log-likelihood under the Student t copula based feature representations is computationally more demanding than computation of the log-likelihood under the Gaussian copula based feature representations, without a substantial improvement in retrieval rate. EMM / ML [1], the multivar. Power Exp. / Geodesic [70] and the proposed Gaussian Copula Weibull / ML approach show similar runtime for the computation of the similarity function but model storage requirements differ drastically. While the multivar. Power Exp. / Geodesic

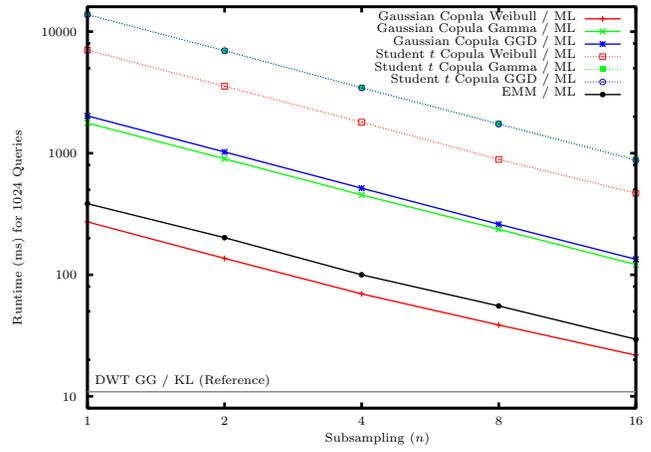


Fig. 6. Runtime reduction due to random (uniform) subsampling to $K = N/n$ query feature vectors (best viewed in color).

approach is quite efficient in terms of storage requirements due to the block-diagonal form of the scatter matrix, the EMM / ML approach requires to store the mixture model component weights, the location vectors as well as the diagonal elements of the covariance matrices. Taking into account the dimensionality of the feature vectors (i.e. 48-dimensional), this leads to the highest storage requirements of our experiments. The Gaussian Copula Weibull / ML approach is ranked in the middle of the field with the same storage requirements as the DTCWT Weibull / KL approach including the additional 18×18 correlation matrix.

As noted in Section III, runtime of the MPE retrieval methods based on the ML selection rule can be improved by reducing the number of query feature vectors by a factor of $1/n$ via uniform random sampling. In Fig. 6, we plot the runtime to perform one image query to 1024 candidates (in milliseconds) on a logarithmic scale as a function of the subsampling factor $n = 1, 2, 4, 8, 16$. As expected, the runtime reduces almost linearly with the number of query feature vectors $K = N/n$. The associated reduction in retrieval rate is depicted in Fig. 5. We highlight that for $n = 4$, the accuracy of the top 16 retrieved images is reduced only by about 2 % points while saving 75 % of the runtime. Retrieval rate of EMM / ML seems to be least affected by subsampling. Note that subsampling reduces the number of query feature vectors, not the storage requirements of the feature representations stated in Table III.

Since the DWT GG / KL approach of Do & Vetterli [5] is commonly-used as a reference in many texture retrieval studies (possibly due to the availability of source code), we include its performance numbers in Fig. 5 and all subsequent figures as a base line as well.

ROC curve plots for the Gaussian Copula Weibull / ML approach including its two major competitors are shown in Fig. 7 for four texture image databases. Note that the ordinate scaling is different for each database. EMM / ML, the Gaussian Copula Weibull / ML and the multivar. Power Exp. / Geodesic schemes score within 5 % points in terms of retrieval rate, but with different rankings depending on the database tested. On the larger STex database (476 images), DWT GG / KL is outperformed by 20 % points over the entire

⁵Available at <http://eigen.tuxfamily.org>.

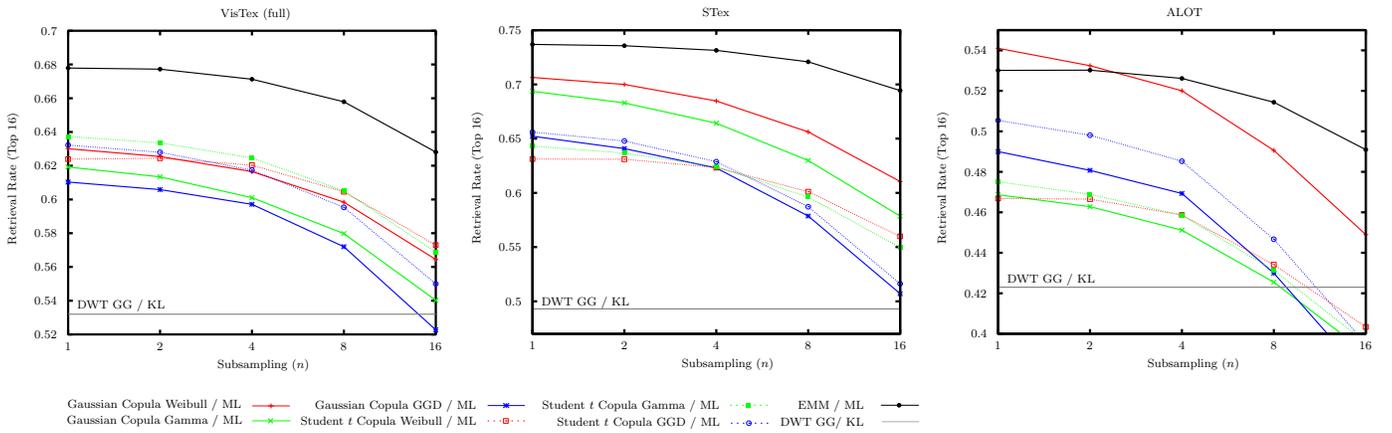


Fig. 5. Retrieval rate reduction due to (uniform) random subsampling to $K = N/n$ query feature vectors $\mathbf{z}_1^*, \dots, \mathbf{z}_K^*$ (best viewed in color).

TABLE III

FEATURE REPRESENTATION (BYTES) AND RUNTIME COMPARISON (IN MS)

Approach	Feature Rep.	Runtime (ms)
DWT GG / KL [5]	432	10.9
Multivar. Power Exp. / Geodesic [70]	432	391.4
Gaussian Copula GG / ML	504	1975.4
Student t Copula GG / ML	512	13800.7
DWT GGamma / KL [16]	648	11.2
DTCWT Weibull / KL [31]	864	18.2
Gaussian Copula Weibull / ML	1656	272.4
Gaussian Copula Gamma / ML	1656	1770.5
Student t Copula Weibull / ML	1664	6948.5
Student t Copula Gamma / ML	1664	13834.5
EMM / ML [2]	6208	384.2

operation range. Based on these results, we conclude that it is imperative to perform extensive tests on several databases to allow quantitative statements about retrieval performance. Testing on just one database might substantially blur the overall impression and lead to false conclusions.

V. DISCUSSION & CONCLUDING REMARKS

Summarizing the results of the experimental section, we conclude that joint statistical modeling of DWT / DTCWT coefficients improves retrieval performance over assuming complete independency among coefficients. We compared several copula based subband models with prior research based on marginal subband models only. Further, we note that overly complex margin models such as the Generalized Gamma distribution do not necessarily improve retrieval rate, possibly due to model parameter estimation issues. This is also true in the context of copula based modeling, since the versatile Student t copula did not show any significant improvements compared to the Gaussian copula.

The results of our study lead to the conclusion that a texture retrieval approach should not only be characterized by its retrieval accuracy but also by its runtime behavior. In this light, the proposed Gaussian copula with Weibull margins approach stands out from the other copula based proposals which exhibit comparable retrieval rates and similar storage requirements but lack competitive runtime behavior.

The first major competitor from the collection of comparative approaches, the EMM / ML approach, has higher storage

requirements for the feature representation and comparable runtime and retrieval performance. However, when reducing the number of query feature vectors, the retrieval rate is more robust at higher subsampling factors (cf. Figs. 5 and 6, $n > 4$). Taking the effort for model parameter estimation into account, we highlight that the EMM / ML approach is based on an EM algorithm (which requires a reasonable initialization), while estimation of the Gaussian copula with Weibull margins only requires to compute a sample correlation matrix and to estimate the Weibull distribution parameters. The latter step can be realized very efficiently by using moment estimates for instance.

The second major competitor, the multivar. Power Exp. / Geodesic approach, is more efficient in terms of storage requirements and exhibits comparable runtime and retrieval performance. However, each new query image requires computation of the feature transformation (i.e. DWT) and feature representation (i.e. estimation of the model parameters). Both, ML and moment estimates are particularly tedious to compute due to the analytically intractable form of the multivariate probability density function. In this context, we like to remind that the ML selection criterion of the MPE framework only requires to compute the feature transformation followed by inserting the query feature vectors into the candidate models of the database. Colloquially speaking, these two retrieval paradigms amount to moving a large amount of data (with possible reduction by subsampling) and evaluating the log-likelihood under a given parametric model versus boiling down the data to a handy feature representation and evaluation of a suitable similarity measure, e.g. KL divergence or geodesic distance.

It is further worth noting why we omitted comparative tests against some of the approaches mentioned in Section I, such as the Refined Histogram technique [17] or a Hidden Markov Tree based approach. First, the Refined Histogram technique is particularly designed to enhance runtime behavior of the GG based margin-only approach of Do & Vetterli [5]. This is accomplished by wedding the modeling results of [19] with the idea of measuring Kullback-Leibler divergences. Since the retrieval rates presented in [17] are only slightly better than the DWT GG / KL approach, we omitted a comparison in

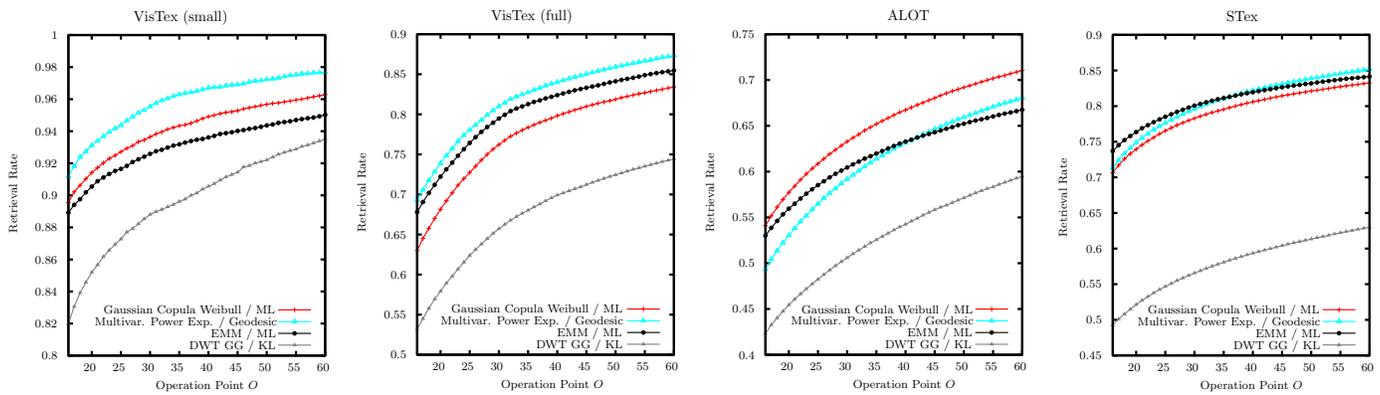


Fig. 7. ROC curve plots for the Gaussian copula with Weibull margins / ML approach vs. the two major competitors, EMM / ML [2] and multivar. Power Exp. / Geodesic [70], as well as the reference approach DWT GG / ML [5] (best viewed in color).

our experimental study. Regarding the family of HMT based retrieval strategies, we point out that we initially performed experiments with the Contourlet HMT approach proposed in [41]. This seems a reasonable choice, since the Contourlet transform is closer to the DTCWT in terms of directional selectivity compared to the DWT. However, it turned out the retrieval improvements, at least on VisTex (small), were minor (i.e. $\sim 2\%$) with respect to [5] and the computational complexity increased to a level comparable with the Student t copula with Gamma margins which performs $\sim 7\%$ better than the baseline (cf. Table II). The disadvantageous runtime behavior can be attributed to the Monte-Carlo approximation of the Kullback-Leibler divergence between two HMTs which we used to measure image similarity (as recommended in [41]). Weighing the considerable runtime demands against the small improvements in retrieval rate, lead to the decision not to include this approach in our experiments either.

As a final remark, we state that the *margin-only* approaches, i.e. DTCWT Weibull / KL, DWT GGamma / KL and DWT GG / KL, excel in efficient computation of the retrieval step due to simple, closed form expressions of the KL divergence, however retrieval rate is clearly inferior to any of the joint statistical models.

REFERENCES

- [1] N. Vasconcelos, "Bayesian models for visual information retrieval," Ph.D. dissertation, Massachusetts Institute of Technology, Massachusetts, United States, 2000.
- [2] N. Vasconcelos and A. Lippman, "A probabilistic architecture for content-based image retrieval," in *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR'00)*, Hilton Head, South Carolina, United States, Jun. 2000, pp. 1216–1221.
- [3] T. Randen and J. Husoy, "Filtering for texture classification: A comparative study," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 4, pp. 291–310, Apr. 1999.
- [4] T. Chang and C. Kuo, "Texture analysis and classification with tree-structured wavelet transform," *IEEE Transactions on Image Processing*, vol. 2, no. 4, pp. 429–441, Oct. 1993.
- [5] M. Do and M. Vetterli, "Wavelet-based texture retrieval using Generalized Gaussian density and Kullback-Leibler distance," *IEEE Transactions on Image Processing*, vol. 11, no. 2, pp. 146–158, Feb. 2002.
- [6] J. Portilla and E. Simoncelli, "A parametric texture model based on joint statistics of complex wavelet coefficients," *International Journal of Computer Vision*, vol. 40, no. 1, pp. 49–70, Oct. 2000.
- [7] G. Tzagkarakis, B. Beferull-Lozano, and P. Tsakalides, "Rotation-invariant texture retrieval with Gaussianized Steerable Pyramids," *IEEE Transactions on Image Processing*, vol. 15, no. 9, pp. 2702–2718, Sep. 2006.
- [8] B. S. Manjunath and W. Y. Ma, "Texture features for browsing and retrieval of image data," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 8, pp. 837–842, Aug. 1996.
- [9] A. Jain and G. G. Healey, "A multiscale representation including opponent color features for texture recognition," *IEEE Transactions on Image Processing*, vol. 7, no. 1, pp. 124–128, Jan. 1998.
- [10] N. Sebe and M. Lew, "Wavelet based texture classification," in *Proceedings of the 15th International Conference on Pattern Recognition (ICPR'00)*, vol. 3, Barcelona, Spain, Sep. 2000, pp. 3959–3962.
- [11] P. Rivaz and N. G. Kingsbury, "Complex wavelet features for fast texture image retrieval," in *Proceedings of the IEEE International Conference on Image Processing (ICIP'99)*, Kobe, Japan, Oct. 1999, pp. 109–113.
- [12] J. Romberg, H. Choi, R. Baraniuk, and N. G. Kingsbury, "Multiscale classification using complex wavelets," in *Proceedings of the IEEE International Conference on Image Processing (ICIP'00)*, vol. 2, Vancouver, Canada, 2000, pp. 371–374.
- [13] C. Shaffrey, N. G. Kingsbury, and I. Jermyn, "Unsupervised image segmentation via markov trees and complex wavelets," in *Proceedings of the IEEE International Conference on Image Processing (ICIP'02)*, vol. 3, Rochester, New York, United States, 2002, pp. 801–804.
- [14] R. J. Clarke, *Transform Coding of Images*. Academic Press, 1985.
- [15] S. Mallat, "Multifrequency channel decompositions of images and wavelet models," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 37, no. 12, pp. 2091–2110, Dec. 1989.
- [16] S. K. Choy and C. S. Tong, "Statistical wavelet subband characterization based on Generalized Gamma density and its application in texture retrieval," *IEEE Transactions on Image Processing*, vol. 19, no. 2, pp. 281–289, Feb. 2010.
- [17] L. Li, C. S. Tong, and S. K. Choy, "Texture classification using refined histogram," *IEEE Transactions on Image Processing*, vol. 19, no. 5, pp. 1371–1378, May 2010.
- [18] M. H. Pi, C. S. Tong, S. K. Choy, and H. Zhang, "A fast and effective model for wavelet subband histograms and its application in texture image retrieval," *IEEE Transactions on Image Processing*, vol. 15, no. 10, pp. 3078–3088, Oct. 2006.
- [19] S. K. Choy and C. S. Tong, "Statistical properties of bit-plane probability model and its application in supervised texture classification," *IEEE Transactions on Image Processing*, vol. 17, no. 8, pp. 1399–1405, Aug. 2008.
- [20] M. J. Wainwright and E. P. Simoncelli, "Scale mixtures of Gaussians and the statistics of natural images," in *Advances in Neural Information Processing Systems (NIPS'99)*, vol. 12. Cambridge, MA: MIT Press, May 2000, pp. 855–861.
- [21] J. Liu and P. Moulin, "Information-theoretic analysis of interscale and intrascale dependencies between image wavelet coefficients," *IEEE Transactions on Image Processing*, vol. 10, no. 11, pp. 1647–1658, Nov. 2001.
- [22] M. Do and M. Vetterli, "Rotation invariant texture characterization and retrieval using steerable wavelet-domain Hidden Markov models," *IEEE Transactions on Multimedia*, vol. 4, no. 4, pp. 517–527, Dec. 2002.

- [23] J. K. Romberg, H. Choi, and R. G. Baraniuk, "Bayesian tree-structured image modeling using wavelet-domain Hidden Markov models," *IEEE Transactions on Image Processing*, vol. 10, no. 7, pp. 1056–1068, Jul. 2001.
- [24] H. Choi and R. G. Baraniuk, "Multiscale image segmentation using wavelet-domain Hidden Markov models," *IEEE Transactions on Image Processing*, vol. 10, no. 9, pp. 1309–1321, Sep. 2001.
- [25] M. S. Crouse, R. D. Nowak, and R. G. Baraniuk, "Wavelet-based statistical signal processing using Hidden Markov models," *IEEE Transactions on Signal Processing, Special Issue on Wavelets and Filterbanks*, vol. 46, no. 2, pp. 886–902, Apr. 1998.
- [26] S. Sakji-Nsibi and A. Benazza-Benyahia, "Indexing of multichannel images in the wavelet transform domain," in *Proceedings of the International Conference on Information and Communication Technologies: From Theory to Practice*, Damascus, Syria, 2008, pp. 1–6.
- [27] —, "Copula-based statistical models for multicomponent image retrieval in the wavelet transform domain," in *IEEE International Conference on Image Processing (ICIP'09)*, Cairo, Egypt, Nov. 2009, pp. 253–256.
- [28] Y. Stitou, N. Lasmar, and Y. Berthoumieu, "Copulas based multivariate Gamma modeling for texture classification," in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP'09)*, Taipei, Taiwan, Apr. 2009, pp. 1045–1048.
- [29] N. G. Kingsbury, "Image processing with complex wavelets," *Phil. Trans. Royal Society London A, a Discussion Meeting on "Wavelets: the key to intermittent information?"*, vol. 357, no. 1760, pp. 2543–2560, Sep. 1999.
- [30] I. W. Selesnick, R. G. Baraniuk, and N. G. Kingsbury, "The dual-tree complex wavelet transform - a coherent framework for multiscale signal and image processing," *IEEE Signal Processing Magazine*, vol. 22, no. 6, pp. 123–151, Nov. 2005.
- [31] R. Kwitt and A. Uhl, "Image similarity measurement by Kullback-Leibler divergences between complex wavelet subband statistics for texture retrieval," in *Proceedings of the IEEE International Conference on Image Processing (ICIP '08)*, San Diego, California, United States, Oct. 2008, pp. 933–936.
- [32] S. Rahman, M. Ahmad, and M. Swamy, "Bayesian wavelet-based image denoising using the Gauss-Hermite expansion," *IEEE Transactions on Image Processing*, vol. 17, no. 10, pp. 1755–1771, Oct. 2008.
- [33] M. A. Miller and N. G. Kingsbury, "Image modeling using interscale phase properties of complex wavelet coefficients," *IEEE Transactions on Image Processing*, vol. 17, no. 9, pp. 1491–1499, Sep. 2008.
- [34] A. Vo and S. Orintara, "A study of relative phase in complex wavelet domain: Property, statistics and applications in texture image retrieval and segmentation," *Signal Processing: Image Communication*, vol. 25, no. 1, pp. 28–46, Jan. 2010.
- [35] R. Kwitt and A. Uhl, "A joint model of complex wavelet coefficients for texture retrieval," in *Proceedings of the IEEE International Conference on Image Processing (ICIP '09)*, Cairo, Egypt, Nov. 2009, pp. 1877–1880.
- [36] —, "Lightweight probabilistic image retrieval," *IEEE Transactions on Image Processing*, vol. 19, no. 1, pp. 241–253, Jan. 2010.
- [37] P. Vandewalle, J. Kovacevic, and M. Vetterli, "Reproducible research in signal processing - What, why, and how," *IEEE Signal Processing Magazine*, vol. 26, no. 3, pp. 37–47, Mar. 2009.
- [38] N. Vasconcelos and A. Lippman, "Library-based coding: A representation for efficient video compression and retrieval," in *Proceedings of the Data Compression Conference (DCC'97)*, Snowbird, Utah, USA, Mar. 1997, pp. 121–130.
- [39] K. Fukunaga, *Introduction to Statistical Pattern Recognition*, 2nd ed. Morgan Kaufmann, 1990.
- [40] N. G. Kingsbury and J. A. Magarey, "Wavelet transforms in image processing," in *Proceedings of the First European Conference on Signal Analysis and Prediction*, Prague, Czech Republic, Jun. 1997, pp. 23–34.
- [41] D. Po and M. N. Do, "Directional multiscale modeling of images using the Contourlet transform," *IEEE Transactions on Image Processing*, vol. 15, no. 6, pp. 1610–1620, Jun. 2006.
- [42] P. S. N. Fisher, "Chi-plots for assessing dependence," *Biometrika*, vol. 72, pp. 253–265, Aug. 1985.
- [43] N. Fisher and P. Switzer, "Graphical assessment of dependence: Is a picture worth 100 tests?" *The American Statistician*, vol. 55, no. 3, pp. 233–239, Aug. 2001.
- [44] B. Everitt, *An R and S-Plus Companion to Multivariate Analysis*. Springer, 2005.
- [45] H. Joe, *Multivariate Models and Dependence Concepts*, ser. Monographs on Statistics and Applied Probability. Chapman & Hall, 1997.
- [46] R. B. Nelsen, *An Introduction to Copulas*, 2nd ed., ser. Springer Series in Statistics. Springer, 2006.
- [47] M. Sklar, "Fonctions de répartition à n dimensions et leurs marges," *Publications de l'Institut de Statistique de l'Université de Paris*, vol. 8, pp. 229–231, 1959.
- [48] M. Rosenblatt, "Remarks on multivariate transformation," *The Annals of Mathematical Statistics*, vol. 23, no. 3, pp. 470–472, 1952.
- [49] M. Varanasi, "Parameter estimation for the Generalized Gaussian noise model," Master's thesis, Rice University, Houston, Texas, 1986.
- [50] K. Krishnamoorthy, *Handbook of Statistical Distributions with Applications*. Chapman & Hall, 2006.
- [51] E. Bouyé, V. Durrelman, A. Nikeghbali, G. Riboulet, and T. Roncalli, "Copulas for finance – a reading guide and some applications," Mar. 2000, working paper.
- [52] C. Genest, B. Rémillard, and D. Beaudoin, "Goodness-of-Fit tests for copulas: A review and a power study," *Mathematics and Economics*, vol. 44, no. 2, pp. 199–213, Apr. 2009.
- [53] D. Berg, "Copula Goodness-of-Fit testing: an overview and power comparison," *The European Journal of Finance*, vol. 15, no. 7, pp. 675–701, Oct. 2009.
- [54] A. K. Nikoloulopoulos and D. Karlis, "Copula model evaluation based on parametric bootstrap," *Computational Statistics and Data Analysis*, vol. 52, pp. 3342–3353, Mar. 2007.
- [55] V. Panchenko, "Goodness-of-Fit tests for copulas," *Physica A*, vol. 355, no. 1, pp. 176–182, Sep. 2005.
- [56] C. Genest and A. C. Favre, "Everything you always wanted to know about Copula modeling and were afraid to ask," *Journal of Hydrological Engineering*, vol. 12, no. 4, pp. 347–368, Jul. 2007.
- [57] M. V. T. Database, "MIT vision and modeling group," [Online], available from: <http://vismod.media.mit.edu/vismod/>.
- [58] H. Akaike, "A new look at the statistical model identification," *IEEE Transactions on Automatic Control*, vol. 19, no. 6, pp. 716–723, Dec. 1974.
- [59] G. Schwarz, "Estimating the dimension of a model," *The Annals of Statistics*, vol. 6, no. 2, pp. 461–464, Mar. 1978.
- [60] C. Genest and B. Rémillard, "Validity of the parametric bootstrap for Goodness-of-Fit testing in semiparametric models," *Annales de l'Institut Henri Poincaré*, vol. 44, no. 6, pp. 1096–1127, Dec. 2008.
- [61] B. Efron and R. Tibshirani, *An Introduction to the Bootstrap*. Chapman & Hall/CRC, 1993.
- [62] A. Genz, "Numerical computation of multivariate normal probabilities," *Journal of Computational and Graphical Statistics*, vol. 1, no. 2, pp. 141–149, Jun. 1992.
- [63] A. Genz and F. Bretz, "Comparison of methods for the computation of multivariate t probabilities," *Journal of Computational and Graphical Statistics*, vol. 11, no. 4, pp. 950–971, Dec. 2002.
- [64] N. A. Heckert and J. J. Filliben, *NIST Handbook 148: DATAPLOT Reference Manual*. National Institute of Standards and Technology Handbook Series, 2003, vol. 1.
- [65] E. de Ves, A. Ruedin, D. Acevedo, C. Benavent, and L. Seijas, "A new wavelet-based texture descriptor for image retrieval," *Lecture Notes in Computer Science, Computer Analysis of Images and Patterns*, vol. 4673, pp. 895–902, Aug. 2007.
- [66] M. Abramowitz and I. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Dover, 1964.
- [67] W. Press and other, *Numerical Recipes in C: The Art of Scientific Computing*, 2nd ed. Cambridge University Press, 1992.
- [68] S. Nadarajah, "A generalized normal distribution," *Journal of Applied Statistics*, vol. 32, pp. 685–694, Sep. 2005.
- [69] G. J. Burghouts and J.-M. Geusebroek, "Material-specific adaptation of color-invariant features," *Pattern Recognition Letters*, vol. 30, no. 3, pp. 306–313, Feb. 2009.
- [70] G. Verdoolaege and P. S. S. De Backer, "Multiscale colour texture retrieval using the geodesic distance between multivariate Generalized Gaussian models," in *Proceedings of the IEEE International Conference on Image Processing (ICIP '08)*, San Diego, California, USA, Oct. 2008, pp. 169–172.
- [71] R. Picard, T. Kabir, and F. Liu, "Real-time recognition with the entire Brodatz texture database," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'93)*, New York, United States, Jun. 1993, pp. 638–639.
- [72] W.-Y. Ma and H. Zhang, "Benchmarking of image features for content-based retrieval," in *Proceedings of the Asilomar Conference on Signals, Systems & Computers*, Pacific Grove, California, United States, 1998, pp. 253–257.

- [73] K. Xu, B. Georgescu, D. Comaniciu, and P. Meer, "Performance analysis in content-based retrieval with textures," in *Proceedings of the International Conference on Pattern Recognition (ICPR'00)*, vol. 4, Barcelona, Spain, Sep. 2000, pp. 4275–4279.
- [74] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society – Series B*, vol. 39, no. 1, pp. 1–38, 1977.
- [75] R. M. Gray, "Vector quantization," *IEEE Transactions on Acoustic Signal and Speech Processing*, vol. 1, pp. 4–29, Apr. 1984.
- [76] E. Gomez, M.-A. Gomez-Vilegas, 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.
- [77] K.-S. Song, "Globally convergent algorithms for estimating Generalized Gamma distributions in fast signal and image processing," *IEEE Transactions on Image Processing*, vol. 17, no. 8, pp. 1233–1250, Aug. 2008.
- [78] N. G. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Applied and Computational Harmonic Analysis*, vol. 10, no. 3, pp. 234–253, May 2001.
- [79] C. J. van Rijsbergen, *Information Retrieval*, 2nd ed. Butterworth-Heinemann, 1979.
- [80] K.-S. Song, "A globally convergent and consistent method for estimating the shape parameter of a Generalized Gaussian distribution," *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 510–527, Feb. 2006.
- [81] T. Ojala, M. Pietikäinen, and D. Harwood, "A comparative study of texture measures with classification based on feature distributions," *Pattern Recognition*, vol. 29, no. 1, pp. 51–59, January 1996.
- [82] E. Hadjidemetriou, M. Grossberg, and S. Nayar, "Multiresolution histograms and their use for recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 7, pp. 831–847, Jul. 2004.
- [83] M. Swain and D. Ballard, "Color indexing," *International Journal of Computer Vision*, vol. 7, no. 1, pp. 11–32, Nov. 1991.



Andreas Uhl is an associate professor at the Department of Computer Sciences (University of Salzburg, Austria) where he leads the Multimedia Processing and Security Lab. His research interests include image and video processing and compression, wavelets, media security, medical imaging, biometrics, and number-theoretical numerics.



Roland Kwitt received his Master and PhD degree in Computer Sciences from the University of Salzburg (Austria) in March 2007 and August 2010. He further holds a Master degree in Telecommunication Systems Engineering from the University of Applied Sciences in Salzburg. He currently pursues post-doctoral studies in LIDAR-based obstacle detection for autonomous driving on railroads. His research interests include statistical image analysis, medical image processing, statistical pattern recognition, image retrieval and multimedia security.



Peter Meerwald received his Master degree in Computer Sciences from Bowling Green State University (Ohio, USA) in 1999 and obtained a Master and PhD degree from the University of Salzburg (Austria) in 2001 and 2010. He is currently pursuing post-doctoral studies at INRIA Rennes (France). His research interests include multimedia security, watermarking, and scalable image and video coding.