Textures and Wavelet-Domain Joint Statistics

Zohreh Azimifar, Paul Fieguth, and Ed Jernigan

Systems Design Engineering, University of Waterloo
Waterloo, Ontario, Canada, N2L 3G1
{szazimif,pfieguth,jernigan}@uwaterloo.ca

Abstract. This paper presents an empirical study of the joint wavelet statistics for textures and other random imagery. There is a growing realization that modeling wavelet coefficients as independent, or at best correlated only across scales, assuming independence within a scale, may be a poor assumption. While recent developments in wavelet-domain Hidden Markov Models (notably HMT-3S) account for within-scale dependencies, we find empirically that wavelet coefficients exhibit within- and across-subband neighborhood activities which are orientation dependent. Surprisingly these structures are not considered by the state-of-the-art wavelet modeling techniques. In this paper we describe possible choices of the wavelet statistical interactions by examining the joint-histograms, correlation coefficients, and the significance of coefficient relationships.

1 Introduction

Statistical models, in particular prior probability models, for underlying textures are of central importance in many image processing applications. However because of the high dimensionality (long-range) of spatial interactions, modeling the statistics of textures is a challenging task. Statistical image modeling can be significantly improved by decomposing the spatial domain pixels into a different basis, most commonly a set of multiscale-multichannel frequency subbands, referred to as the wavelet domain [1]. Indeed, the wavelet transform (WT) has widely been used as an approximate whitener of statistical time series. It has, however, long been recognized [2] that the wavelet coefficients are neither Gaussian, in terms of the marginal statistics, nor white, in terms of the joint statistics.

The wavelet parsimony representation observes that the majority of the coefficients happen to be small, and only a few of the coefficients are large in magnitude, implying that the marginal distributions of the high frequency wavelet subbands are more heavily tailed than a Gaussian, with a large peak at zero. Existing works assume a generalized Gaussian model, some sort of mixture, for the marginal distribution [1]. Chipman et al. [1] and Crouse et al. [2] showed that this heavy-tailed non-Gaussian marginal can be well approximated by a Gaussian Mixture Model (GMM). Accordingly, wavelet non-linear shrinkage, such as Bayesian estimation has been achieved with these non-Gaussian priors, which consider this kurtosis behavior of the wavelet coefficients.

A. Campilho, M. Kamel (Eds.): ICIAR 2004, LNCS 3212, pp. 331–338, 2004.
© Springer-Verlag Berlin Heidelberg 2004

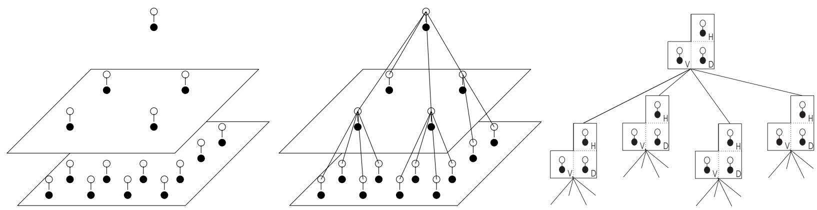

Fig. 1. Three hidden Markov models. The empty circles show the hidden states and black nodes are the coefficient values. (a) Hidden states are independent. (b) Interscale dependencies are modeled. (c) The three subbands are integrated into one hybrid HMT.

A opposed to the marginal models, the question of joint models is much more complicated and admits for more possibilities, with structures possible across subbands, orientations, and scales. Since the development of zerotree coding for image compression there have been many efforts to model these structures, including Markov random fields (MRFs) [3], [4], Besov spaces [5], and the wavelet hidden Markov models (HMMs) [2], [6], [7] and Gaussian scale mixture(GSM) [8]. The wavelet-based HMMs, in particular, have been thoroughly studied and successfully outperform many wavelet-based techniques in Bayesian denoising, estimation, texture analysis, synthesis and segmentation.

HMMs are indeed intended to characterize the wavelet joint statistics. As visualized by Fig. 1, the class of the HMMs mainly includes Independent Mixture Model (IMM) [2], Hidden Markov Tree (HMT) [2], and HMM-3S [7]. In general, they adopt a probabilistic graph, in which every wavelet coefficient (node) is associated with a set of discrete hidden states $S = 0, 1, \ldots, M-1$ (in particular $M = 2$) displayed as empty circles in Fig. 1. To model the connectivity of those states, HMMs first define some hypothesis based on the wavelet coefficients properties, then parameterize models that fit into those assumptions and can be solved by existing algorithms.

In the two-state IMM, the simplest case of HMMs, hidden states are assumed to be independent and every wavelet coefficient is modeled as Gaussian, given its hidden state value (the variance). More sophisticated approaches sought to model the local wavelet statistics by introducing Markovian dependencies between the hidden state variables across scales and orientations. Crouse et al. [2] introduced the HMT, which captures wavelet interscale dependencies by considering Markov chains across scales, while assuming independence within and across the three high frequency channels. Fan and Xia [7] proposed HMT-3S in which, in addition to the joint interscale statistics captured by HMT, the dependencies across subbands are exploited by integrating three corresponding coefficients across three orientations.

Goal of this paper: Motivated by these inter-coefficient probabilistic studies, the primary goal of this paper is to study the wavelet joint statistics by empirically investigating local random field neighborhoods representing statistics of within- and across-scale coefficients. Although the previous observations

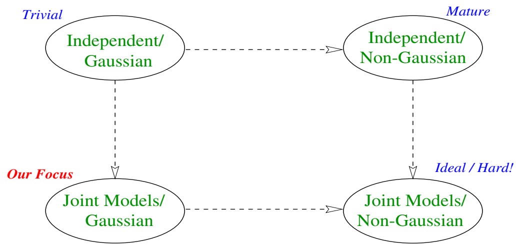

Fig. 2. Focus of this work: Development of joint Gaussian models of wavelet statistics.

highlighted some main wavelet coefficient correlation, there is still uncertainty in wavelet statistics as to whether these approaches offer reasonable choices of correlations? How should one examine sufficiency of wavelet models? Would these models be justified by empirical statistics? This paper is meant to discuss these issues and demonstrate the structure of coefficient correlation that was not captured by HMMs due to their primary assumptions.

Development of wavelet Gaussian random field models for statistical textures forms the focus of our work (shown in Fig. 2). The goal, of course, is the development of non-Gaussian joint models with non-trivial neighborhood. However for the purpose of this paper, we are willing to limit ourselves to simplifying marginal assumptions (Gaussianity) which we know to be incorrect, but which allow us to undertake a correspondingly more sophisticated study of joint models.

Example joint histograms as representatives of the underlying coefficients densities are visualized. We display the hierarchy of wavelet covariance structure and define statistical neighborhoods for the coefficients. The main novelty is the systematic approach we have taken to study the wavelet neighborhood system including 1) inter-scale dependency, 2) within-scale clustering, and 3) across-orientation (geometrical constraints) activities. This probabilistic modeling is directly applied to the wavelet coefficient values, but to some extent their significance is also considered. Surprisingly our empirical observation indicates that the wavelet correlation structure for different textures does not always match with those offered by the HMMs. We will discuss this in later sections.

2 Wavelet Neighborhood Modeling

In order to study exact correlations between the wavelet coefficients we considered a class of statistical textures based on Gaussian Markov random field (GMRF) covariance structures, as shown in Fig. 3. They are spatially stationary, an assumption for convenience only and is not fundamental to our analysis.

The chosen spatial domain covariance structure P_s is projected into the wavelet domain by computing the 2-D WT W, containing all translated and dilated versions of the selected wavelet basis functions:

$$P_w = W P_s W^T \tag{1}$$

where we have restricted our attention to the set of Daubechies basis functions.

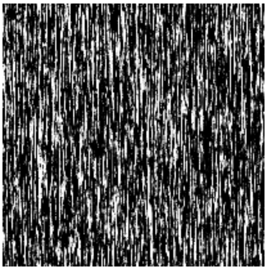

(a) Tree-bark

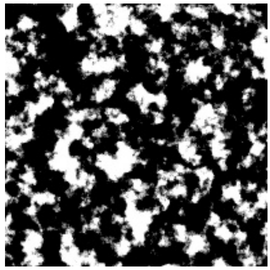

(b) Thin-plate

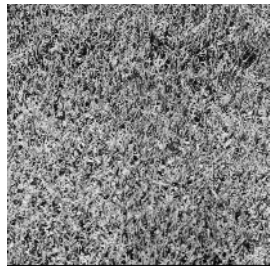

(c) Grass

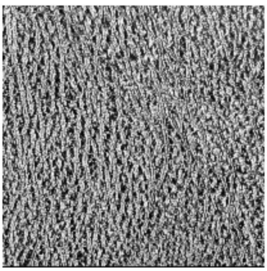

(d) Calf leather

(e) Pigskin

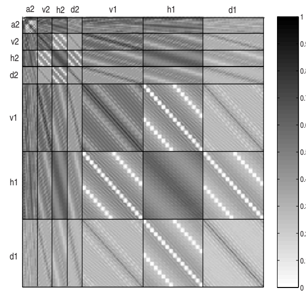

(f) P_w

Fig. 3. (a-e) Five GMRF textures used to visualized wavelet correlation structure. (f) Correlation coefficients of a spatial thin-plate model in the wavelet domain, in which the main diagonal blocks correspond to the same scale and orientation, whereas off-diagonal blocks illustrate cross-correlations across orientations or across scales.

The wavelet covariance, P_w (Fig. 3(f)), is not a diagonal matrix, indicating that the wavelet coefficients are not independent. Intuitively, localized image structures such as edges tend to have substantial power across many scales. More interestingly, P_w is block-structured, and it is evident that the coefficients interactions align with direction of their subband. We have observed [9] that, although the majority of correlations are very close to zero (i.e., decorrelated), a relatively significant percentage (10%) of the coefficients are strongly correlated across several scales or within a particular scale but across three orientation subbands. Clearly a random field model for wavelet coefficients will need to be explicitly hierarchical. One approach to statistically model these relationships was to implement a multiscale model [9]. Although the multiscale model captured the existing strong parent-child correlation, spatial and inter-orientation interactions are not explicitly taken into consideration. Our most recent work [10] investigated two techniques to approximate non-Markov structure of P_w into a Markovian neighborhood which contains the significance of inter-orientation and spatial relationships, which we seek to visualize more formally and compare with other methods in this paper.

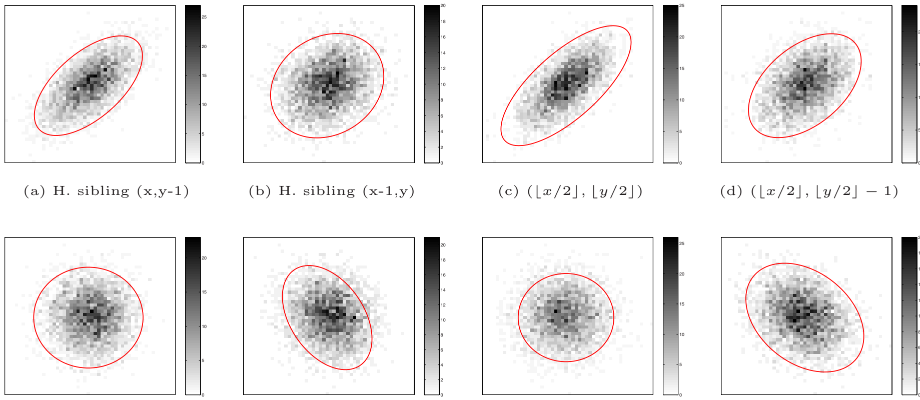

(a) H. sibling (x,y-1) (b) H. sibling (x-1,y) (c) $(\lfloor x/2 \rfloor, \lfloor y/2 \rfloor)$ (d) $(\lfloor x/2 \rfloor, \lfloor y/2 \rfloor - 1)$

(e) V. cousin (x,y) (f) V. cousin (x-1,y-1) (g) D. cousin (x,y) (h) D. cousin (x,y-1)

Fig. 4. Empirical joint histograms of a coefficient at position (x,y) in a horizontal subband associated with different pairs of coefficients at the same scale and orientation (a,b), at the same orientation but adjacent scales (c,d), at the same scale but across orientations (e-h). The skewness in the ellipsoid indicates correlation.

2.1 Wavelet Domain Joint Histograms

In order to characterize the wavelet neighborhood explicitly, we first utilize joint histogram plots. This intermediate step helps to identify two coefficients' dependency even if they show as decorrelated on their correlation map (i.e. decorrelation does not always mean independence!).

For a typical texture, joint histograms of a horizontally aligned coefficient at position (x,y) associated with different pairs of coefficients are illustrated in Fig. 4. These plots highlight the following important aspects of the coefficients connectivity:

Remark 1: In the top row, the first two plots show extended contours indicating that two spatially adjacent horizontal coefficients not only are dependent but also the direction of their correlation matches with that of their subband. For instance, within its subband, a horizontal coefficient is more correlated with its adjacent left and right neighbors than up and down neighbors.

Remark 2: The top row's last two plots are joint histograms of parent-child horizontal coefficients. It is a quite evident that a child strongly depends not only on its parent (a fact observed by many other researchers) but also on its parent's adjacent neighbor (left or right). We also observed that, by symmetric, a vertical coefficient statistically depends on its parent and parent's upper or lower neighbor.

Remark 3: The bottom row plots display joint histograms of a horizontal coefficient with its corresponding neighbors within the same scale but across other two orientations. Firstly, the nearly circular contours indicate that coefficients at the same location but from different orientations are almost independent! Secondly, there is still some inter-orientation correlation which aligns with direction of the centered coefficient (i.e. correlation structure is subband dependent).

In summary, we emphasize that this paper is not to report the striking wavelet correlations exhibited in these empirical observations. Rather, it is observed that, surprisingly, the existing wavelet joint models not only consider a *subset* of these inter-relationships but also fail in connecting some coefficients which are indeed independent, e.g. in HMT-3S three coefficients at the same location from three subbands are grouped into one node (assumed to be correlated), an assumption that is rejected by these histogram plots.

2.2 Wavelet Domain Correlation Structure

Being motivated by the histogram plots, we have chosen to study the problem visually, and without any particular assumption regarding the coefficient position on the wavelet tree. First, correlation coefficients are calculated from the wavelet prior P_w for three fine scale textures displayed in Fig. 3. As shown in Fig. 5, we use the traditional 2-D wavelet plot to display the correlation of a coefficient coupled with any other coefficient on the entire wavelet tree. Each panel includes three plots illustrating local neighborhood for a centered coefficient (marked by •) chosen from horizontal, vertical, and diagonal subbands. The left column panels in Fig. 5(a-c), show correlation coefficients for a coefficient paired with all other nodes on the wavelet tree.

There is a clear consistency between the joint histograms and these correlation maps which shows 1) The concentration of the wavelet correlations in a locality. 2) This locality increases toward finer scales, which supports the persistency property of wavelet coefficients [6]. 3) The local neighborhood definition for any given pixel is not limited to the pixel's subband: it extends to dependencies across directions and resolutions. Besides the long range across scale correlations, every typical coefficient exhibits strong correlation with its spatially near neighbors both within subband and across orientations. 4) The correlation structure for horizontally and vertically aligned coefficients are almost symmetrically identical. For textures whose edges extend more or less toward one direction (such as tree-bark), this similarity does not hold.

To consider the sparse representation property of the WT, these empirical evaluations have been extended to dependency structure of those significant coefficients. In [9], we defined the significance map as a tool to identify those correlations corresponding to the significant coefficients. Fig. 5(d-f) show the significance of correlations for the corresponding panels displayed in Fig. 5(a-c). It is evident from these diagrams that within scale dependency range reduces to shorter locality (yet orientation dependent), but across scale activities still present up to several scales.

Interestingly, the wavelet correlation plots in Fig. 5 show a clear consistency in structure for many textures. They confirm that 1) The well-structured coefficients dependencies are hierarchical and orientation dependent. 2) Coefficients across three orientations and at the same spatial position are decorrelated, however, there is a clear dependency between coefficients across orientations and

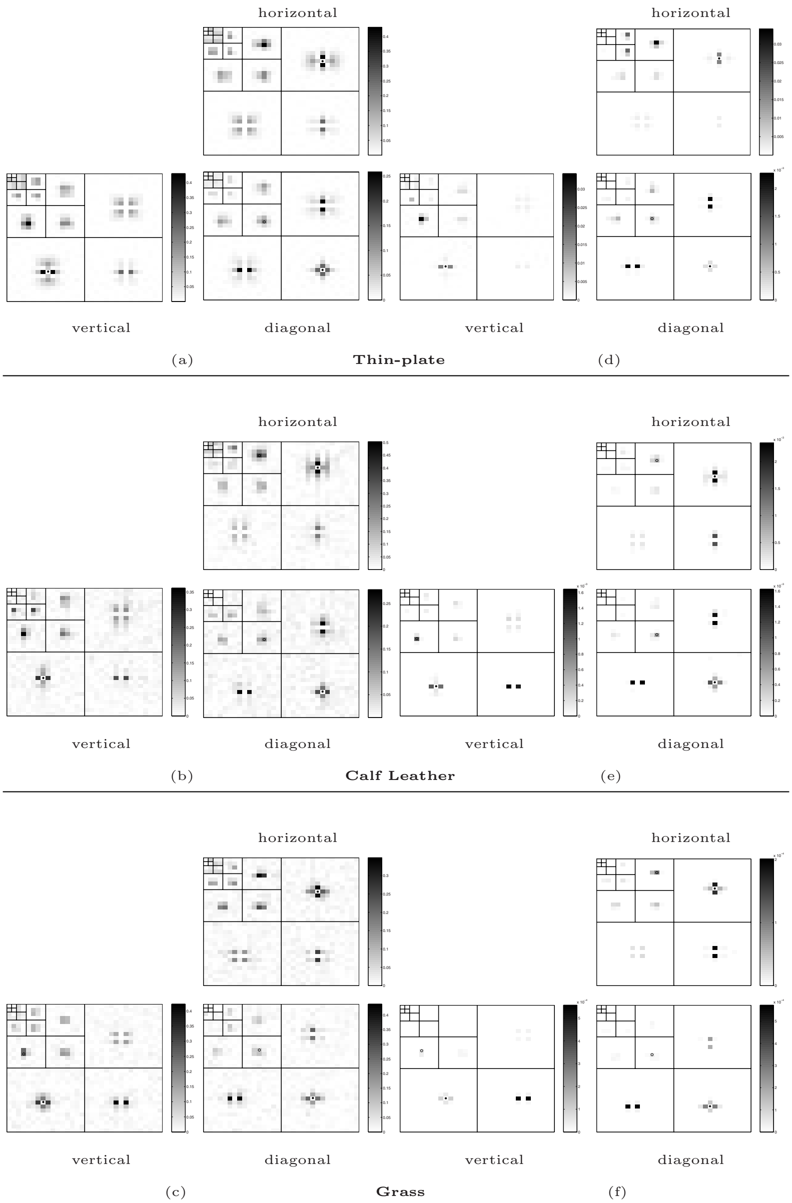

Fig. 5. Wavelet correlation structure for three fine scale textures displayed in Fig. 3. Each panel contains three plots illustrating local neighborhood for a centered coefficient (marked by •) from horizontal, vertical, and diagonal subbands. The left column panels (a-c) show correlation coefficients for a coefficient paired with all other nodes on the wavelet tree. The right column (d-f) are plots of significance of above interrelationships.

at nearby neighborhood. 3) Coefficients are correlated with their parent and neighbors of it, in addition to parents across other two coarser subbands.

3 Conclusions

A thorough study of the 2-D wavelet statistics has been presented in this paper. Empirical examination of the coefficient correlations, within or across scales, revealed the fact the there exist local and sparse random field models governing these local dependencies.

A superset including all statistically local neighbors for a wavelet coefficient was demonstrated. We compared our modeling observations with the advanced wavelet joint models. This study showed that the correlation structures presumed and proposed by those approaches (such as HMT-3S) does not always accurately integrate the correlated coefficients. We also discussed examples of interscale and intra-scale dependencies that are missing in the existing models. We are expanding this ongoing research to the statistics of real world images. The early empirical examinations show consistency with the correlation structures studied in this article.

References

1. H. Chipman, E. Kolaczyk and R. McCulloch, “Adaptive Bayesian wavelet shrinkage”, J. Amer. Statis. Assoc., pp. 92-99, 1997
2. M. S. Crouse, R. D. Nowak, R. G. Baraniuk, “Wavelet-based statistical signal processing using hidden Markov models”, IEEE trans. on SP, vol. 46, pp. 886-902, 1998
3. M. Malfait and D. Roose, “Wavelet-based image denoising using a Markov random field a priori model”, IEEE Trans. on IP, vol. 6, pp. 549-565, 1997
4. A. Pizurica, W. Philips, I. Lemahieu, and M. Acheroy, “A joint inter- and intrascale statistical model for Bayesian wavelet based image denoising”, IEEE Trans. on IP, vol. 11, pp. 545-557, 2002
5. A. Srivastava, “Stochastic models for capturing image variability”, IEEE Signal Processing Magazine, vol. 19, pp. 63-76, 2002
6. J. Romberg, H. Choi, and R. Baraniuk, “Bayesian tree-structured image modeling using wavelet-domain hidden Markov models”, IEEE trans. on IP, vol. 10, pp. 1056-1068, 2001
7. G. Fan and X. Xia, “Wavelet-based texture analysis and synthesis using hidden Markov models”, IEEE Trans. on Cir. and Sys., vol. 50, pp. 106-120, 2003
8. J. Portilla, V. Strela, M. J. Wainwright and E.P. Simoncelli, “Image denoising using Gaussian scale mixtures in the wavelet domain”, IEEE Trans. on IP, vol. 12, pp. 1338-1351, 2003
9. Z. Azimifar, P. Fieguth, and E. Jernigan. “Towards random field modeling of wavelet statistics.” Proceedings of the 9th ICIP, 2002.
10. Z. Azimifar, P. Fieguth, and E. Jernigan. “Hierarchical Markov models for wavelet-domain statistics.” Proceedings of the 12th IEEE SSP, 2003.