

A Detection-Theoretic Approach to Texture and Edge Discrimination

Iasonas Kokkinos and Petros Maragos

Computer Vision, Speech Communication and Signal Processing Group

School of Electrical and Computer Engineering

National Technical University of Athens, Greece

jkokkin@cs.ntua.gr, maragos@cs.ntua.gr*

Abstract

We present a probabilistic approach to the discrimination between textured areas and edges; locally defined probabilistic models are used, which model textured areas as sinusoidal and edges as phase-congruent signals. We build a link with energy-based feature detection and propose a simple approach to the discrimination between these two classes. This facilitates the selective use of texture/edge features for edge detection and image segmentation applications. Our experimental results demonstrate the applicability of our method on natural images.

1. Introduction

Two of the best studied computer vision problems are texture analysis and edge detection; their applications to a broad range of vision areas make them fundamental elements of any front-end vision system. A standard processing stage for both problems is filtering with band-pass filterbanks, which is inspired by the structure of the human visual system and aims at providing a representation of the input image that is better suited for these tasks than the raw intensity values. Specifically, for edge detection one expects the filtered signal to indicate potential object borders, while for texture analysis one wishes to derive well behaved signals, capturing information about local image structure while lending themselves to other tasks, like shape-from-texture or segmentation.

The filters used for these two distinct tasks have similar frequency response characteristics, or are even identical and can easily confuse textured areas with edges; this results on the one hand in false edge positives on textured areas and strong texture features along edges. Some recent approaches related to this problem include [5, 13], where it

is attempted to automatically determine which areas in the image are textured. Herein we extend our previous work in [9] and present a common framework for modeling texture, edge and smooth areas that deals with the problem in a simple and efficient manner. We use a set of models that quantize the possible appearance variations of an image neighborhood and use a probabilistic criterion to decide which of these best describes the image in this neighborhood. Formulating this procedure as an M-ary hypothesis detection problem [18, 7] results in a soft assignment of a texture/edge/smooth tag to each image pixel.

The discrimination between texture and edges is evaluated on a small edge detection task using the Berkeley image database; we have also integrated our results with the Dominant Components Analysis (DCA) [6] algorithm for texture analysis, that entails a related channel selection procedure. Segmentation experiments using the fused features are shown, giving promising results.

2. Quadrature Filter Pairs and Local Energy-Based Feature Detection

A quadrature pair of filters consists of one even- and one odd-symmetric filter, h_e, h_o , that have zero mean and have identical Fourier spectrum magnitude. The use of QFPs is a popular method in edge detection applications for simultaneously detecting and localizing both step- and line-like edges [15, 16, 1]. Their application relies on the observation that we perceive edges at areas which can be locally modelled as a sum of *phase congruent* [15] sinusoids:

$$O(x) \simeq I_E(x) = A \sum_k a_k \cos(\omega_0 kx + \phi) + B. \quad (1)$$

Above and in all of the following, we use the following notation: O is the observed signal, I_M is its approximation using the model M (in the case above, E for edges), ω_0 is the fundamental frequency of the Fourier series, ϕ is a phase offset which is common for all harmonic components and B is a constant term (DC). The approximation holds

*This work was partially supported by the Greek research program HRAKLEITOS, which is co-funded by the European Social Fund (75%) and National Resources (25%), the European NoE MUSCLE and the European STREP HIWIRE

around point $x = 0$, where the feature related to the model is perceived. Since we do not deal with 2-D features like junctions and corners, the following analysis uses the 1-D profiles of the features along the orientation of their variation.

The model of Eqn. (1) captures both line- and step- like edges ($\phi = 0, \pi/2$ respectively) and we shall use it in what follows for edges. QFPs fit ideally with it, since they can provide a *phase invariant* quantity: the *local energy* measurement $LE(x) = (h_e * O)^2 + (h_o * O)^2$, estimated using any QFP will not depend on ϕ , so this quantity will respond in the same manner to all signals of the form of Eqn. (1). We note that the signal defined by Eqn. (1) is periodic, but the spatial support of the QFPs is chosen to be small enough to take into account less than one half of the signal's period, where it behaves like a typical edge signal; examples of QFPs used for edge detection can be found in [15, 16, 1].

The odd and even parts of complex Gabor wavelets, which are ubiquitous in texture analysis algorithms [2, 3, 6] act like a QFP, with the sole difference that the even filter has a non-zero DC component; references to methods facing this problem can be found in [1]. In a manner analogous to edge detection, the energy of this QFP offers a phase-invariant quantity that can be used as a local measure of the signal spectrum energy around the peak frequency of the Gabor filters. A simple model for texture (T) that we have shown in [9] to underlie the use of Gabor filters is that of a locally sinusoidal signal:

$$O(x) \simeq I_T(x) = A \cos(\omega_0 x + \phi) + B. \quad (2)$$

This is among the simplest and most tractable texture models; it lies at the heart of more complicated ones, like the multicomponent AM-FM models [6] that can accurately reproduce complex texture patterns. Contrary to Eqn. (1), Eqn. (2) models signals that are better localized in frequency, i.e. smoothly varying sinusoidals; another difference is that the support of the related Gabor QFP typically allows for several oscillations of the harmonic component, thereby rendering their outputs better tuned to periodic signals than isolated intensity variations.

Smooth (S) regions are modelled as constant signals:

$$O(x) \simeq I_S(x) = B. \quad (3)$$

This model complements the previous two and reduces the credibility of texture features at smooth image areas.

3 Detection Theoretic Formulation of Local Energy Feature Detection

In what follows we assume that images can be *locally* described adequately in terms of one of the models of Eqns. (1)-(3), and derive expressions for their probabilities using local energies estimated by corresponding QFPs.

3.1 Generative Model

Our approach is based on the use of generative models, i.e. probabilistic models with a small set of parameters which can be used to reproduce the image observations: each model M uses a K dimensional basis \mathcal{B}_M to approximate the appearance of the observations O around the point $x = 0$, giving the model-based prediction $I_M(x; \mathcal{A}_M) = \sum_{i=1}^K A_i \mathcal{B}_{M,i}(x)$; the expansion coefficients $\mathcal{A}_M = \{A_1, \dots, A_K\}$ constitute the parameter set of the model, while differences between the synthesized and observed signals are accounted for by a noise process. When a set of hypotheses $\mathcal{H} = \{H_1, H_2, \dots, H_N\}$ compete to explain the same observations in terms of alternative generative models, the decision among one of them can be formulated as a Generalized Likelihood Ratio Test (GLRT) decision process [18, 7]: the ML parameter set estimate $\hat{\mathcal{A}}_i$ for each hypothesis i is inserted into the likelihood expressions, giving rise to posterior expressions using Bayes' rule:

$$P(H_i|O) = \frac{P(O|\hat{\mathcal{A}}_i, H_i)}{\sum_{H_j \in \mathcal{H}} P(O|\hat{\mathcal{A}}_j, H_j)}. \quad (4)$$

Probability expressions related to the models of Eqns. (1)-(3) can be derived using the generative model introduced in [9], where it was used to label image pixels as being either textured or smooth; herein we show that this approach carries over directly to edge features. The model we use incorporates the locality of the decision process using a confidence value, $G(x)$ associated with the predictions at point x . The quantity $G(x)$ decreases with the distance from the point $x = 0$, and a background hypothesis is introduced to account for the appearance variation away from it. This can be formalized by the introduction of a binary random variable z_x which indicates whether the observation at point x is due to the fore- or background hypothesis; what we previously described amounts to setting $P(z_x = 1|x) = G(x)$. Assuming independent noise and taking the background model to be a uniform distribution, the likelihood of the observation $O(x)$ at point x under hypothesis i writes:

$$\begin{aligned} P(O(x)|x, i) &= \sum_{z_x \in \{0,1\}} P(O(x), z_x|x, i) \\ &= \sum_{z_x \in \{0,1\}} P(O(x)|z_x, x, i) P(z_x|x, i) \\ &= \underbrace{G(x) \int_{\mathcal{A}} P_i(O(x)|I_i(x; \mathcal{A}_i)) dP(\mathcal{A}_i)}_{z_x=1} + \underbrace{(1 - G(x))c}_{z_x=0}. \end{aligned} \quad (5)$$

P_i expresses how likely observation $O(x)$ is given the synthesis of hypothesis i , $I_i(x; \mathcal{A}_i)$ at point x , while c is a constant likelihood term assigned to any observation from

the background distribution. We have used $P(z_x|x, i) = P(z_x|x) = G(x)$ since we assume that the spatial decay of our confidence in the hypothesis prediction is common for all three models. Using the independence of the noise process we can write:

$$\log P(O|i) = \sum_x \log P(O(x)|x, i). \quad (6)$$

To obtain tractable expressions we bring the summation used for $P(O(x)|x, i)$ in Eqn. (5) outside the logarithm; we use the concavity of the log function and apply Jensen's inequality to derive the following lower bound on log P :

$$LB(O|i) = \sum_x G(x) \log P_i(O(x)|x, i) + \sum_x (1-G(x)) \log c,$$

where for brevity we have replaced the integral in Eqn. (5) with the expression $P_i(O(x)|x, i)$. In what follows we shall use this lower bound instead of the original expression for the data likelihood; even though it is not guaranteed to behave in the same manner as log P , it is indicative of its behavior. Another approximation is that we consider the integral over \mathcal{A} as approximately equal to the value of the integrated function at the ML estimate, $\hat{\mathcal{A}}$:

$$\sum_x G(x) \log P_i(O(x)|x, i) \simeq \sum_x G(x) \log P_i(O(x)|I_i(x; \hat{\mathcal{A}}_i))$$

where $\hat{\mathcal{A}}$ is now chosen to maximize the lower bound instead of the original likelihood expression.

This model is different from those used by other authors [12, 5], in that there is a built-in 'mixture modeling' aspect in Eqn. (5) accounting for the locality of the modeling process. In [12, 5] all the reconstruction errors are given equal weight, so $LB(O|i) = \sum_x \log P(O(x)|x, i)$; the goal in that case is to accurately reconstruct the whole image using local features, while in our case it is to model the image in the neighborhood around point 0.

In the case of independent Gaussian noise with deviation σ , the expression for $\log P_i(O(x)|I_i(x; \mathcal{A}_i))$ becomes $-\frac{(I_i(x; \mathcal{A}_i) - O(x))^2}{2\sigma^2} - \log \sqrt{2\pi}\sigma$ so that we can write for a fixed set of parameters \mathcal{A}

$$LB(O|i; \mathcal{A}) = - \sum_x G(x) \frac{(I_i(x; \mathcal{A}) - O(x))^2}{2\sigma^2} - \sum_x G(x) \log \sqrt{2\pi}\sigma + c', \quad (7)$$

where $c' = \sum_x (1 - G(x)) \log c$. Given that the confidence function $G(x)$ has the same shape for all three models, the lower bound expressions are commensurate for different hypotheses.

Our feature models fit naturally in this framework; for example expanding Eqn. (1) gives

$$I(x; \mathcal{A}_E) = A_e \mathcal{B}_{E,e}(x) + A_o \mathcal{B}_{E,o}(x) + A_d \mathcal{B}_{E,d}(x) \quad (8)$$

which expresses the synthesis of the edge model as an expansion on a linear basis formed from an even, an odd and a DC signal; in the expression above $A_e = A \cos(\phi)$, $A_o = -A \sin(\phi)$, $A_d = B$ and the basis elements are the signals $\mathcal{B}_{E,e}(x) = \sum_{k=1}^N a_k \cos(\omega_0 kx)$, $\mathcal{B}_{E,o}(x) = \sum_{k=1}^N a_k \sin(\omega_0 kx)$ and $\mathcal{B}_{E,d}(x) = 1$. Similarly we get a 3-D basis for the sinusoidal case while for the smooth case we have a 1-D basis; apart from the spatially varying confidence measure $G(x)$ this is the typical linear model used in detection theory [7]. In the following we drop the model index from the basis expressions, unless necessary, since the treatment is common for all three models.

3.2 Local Energy & Observation Likelihood

Using Eqn. (8) the maximization of Eqn. (7) with respect to $\{A, \phi, B\}$ can be interpreted as the *weighted least squares projection* of O onto a linear basis $\mathcal{B}_e(x), \mathcal{B}_o(x), \mathcal{B}_d(x)$, with weights equal to $G(x)$; A, ϕ can then be expressed as $A = \sqrt{A_e^2 + A_o^2}$, $\phi = \tan^{-1} \frac{A_e}{A_o}$.

For the case where $\sum_x G(x) \mathcal{B}_e(x) = 0$, the maximum likelihood condition for each expansion coefficient A_b does not involve the rest and we have:

$$\left. \frac{\partial LB}{\partial A_b} \right|_{A_b = \hat{A}_b} = 0 \rightarrow \hat{A}_b = \frac{G \mathcal{B}_b \cdot O}{G \mathcal{B}_b \cdot \mathcal{B}_b}. \quad (9)$$

The notation $G \mathcal{B} \cdot O$ stands for $\sum_x G(x) \mathcal{B}(x) O(x)$; The estimation of \hat{A}_b can thus be implemented using filtering with $G(x) \mathcal{B}_b(x)$, followed by normalization; in general we can show that filtering operations with QFPs can be used to derive the parameter set $\hat{\mathcal{A}}$ that maximizes Eqn. (7) for each of the three models. Specifically, when $G \mathcal{B}_i \cdot \mathcal{B}_j \neq 0$ for some $i \neq j$, the estimates can be derived by adopting to our case the technique of normalized differential convolution [8, 19], which is used for projecting data with unknown certainty on bases with varying applicability. In our scenario the basis elements are the even/odd and DC signals of the generative models, the data certainty is equal to 1 wherever we have an observation, and the basis applicability is equal to $G(x)$. The equivalence with weighted least squares projection is detailed in [19].

Before proceeding, we clarify what these expressions imply, using the texture model as example: $I_T(x) = A \cos(\omega_0 x + \phi) + B$. In this case $\mathcal{B}_{T,e}$ and $\mathcal{B}_{T,o}$ constitute a sine/cosine pair, and modeling $O(x)$ in terms of $I_T(x)$ amounts to estimating the values of A, ϕ for which Eqn. (2) best fits $O(x)$ locally. Assuming G has the form of a Gaussian and ignoring any DC term, estimating \hat{A}_e and \hat{A}_o as in Eqn. (9) amounts to convolving the signal with even and odd Gabor filters. We can thus interpret convolving with a Gabor filter as calculating *the optimal weighted projection of an image neighborhood onto a sinusoidal basis*. Along the same line, we can view filtering with QFPs

for edge detection as estimating the optimal weighted projection of the image neighborhood onto the edge-function basis, $\mathcal{B}_{E,e}, \mathcal{B}_{E,o}$; at last, the output of convolving with a Gaussian function can be seen as an optimal weighted projection of the observed image data on the basis element 1.

If $\sum_x G(x)\mathcal{B}_e(x) = 0$, one can show that:

$$\begin{aligned} & - \sum_x G(x) \left(O(x) - \sum_b \hat{A}_b \mathcal{B}_b \right)^2 = \\ & - \sum_x G(x) O(x)^2 + \sum_b \hat{A}_b^2 \sum_x G(x) \mathcal{B}_b(x)^2. \end{aligned}$$

Normalizing the basis elements so that $\sum_x G(x)\mathcal{B}_b(x)^2 = \sum_x G(x)$ and considering all other terms in Eqn. (7) as constant, the lower bound can be written as a function linearly increasing with $\sum_b \hat{A}_b^2$:

$$LB = \frac{\sum_x G(x)}{2\sigma^2} \sum_b \hat{A}_b^2 + c''. \quad (10)$$

The quantity c'' is determined by $\sum_x G(x)I(x)^2$, $\sum_x G(x)$, σ and c' , which are the same for all three models. For both the texture and edge models, in the case of zero mean signals the quantity $LB_i = \sum_x G(x)/2\sigma^2 \sum_{b \in \{e,o\}} A_{i,b}^2$, $i \in \{T, E\}$ is a scaled version of the local energy derived from filtering with the QFPs used for estimating the ML estimates. For the sinusoidal case these are Gabor filters and for the edge detection case these are equal with the basis elements multiplied by $G(x)$. This establishes the link between the local energy measurements and the log-likelihood of the data using the generative models presented in the beginning of this section.

Using the expression in Eqn. (10) instead of the log-likelihoods the posterior probability of hypothesis i writes:

$$P(i|O) = \frac{\exp(c_M \sum_b \hat{A}_{i,b}^2)}{\sum_{j \in E,T,S} \exp(c_M \sum_b \hat{A}_{j,b}^2)}, \quad c_M = \frac{\sum_x G(x)}{2\sigma^2}$$

An intricacy concerns the smooth hypothesis model, which uses less parameters and is therefore bound to perform worse than the other two models. We therefore introduce an MDL-like term in Eqn. (10), which favors it in the absence of strong variations in the image, so $\exp(LB_S) = c_{MDL} \exp(c_M A_{S,d}^2)$; c_{MDL} is allowed to vary according to how easily we want to label a pixel as non-smooth.

3.3 Multiple Scales

Up to now we have implicitly assumed that all three models attempt to explain a neighborhood of the image at a fixed scale; as we pass from one scale to another the decision made may change. For example the neighborhood of a

point on the crest of a slowly varying sinusoid may be considered smooth at a fine scale, and as texture at a larger one; at the large scale the data have a larger amount of energy, which is explained by the texture hypothesis, while at the small scale the MDL criterion forces a simpler explanation.

A problem that emerges is that the likelihood terms include a summation over a varying number of observations, so the model-based likelihood terms as well as the residual quantities c'' render the expressions at multiple scales incommensurate. We therefore apply the scale normalization principle of [11] and derive an expression that is invariant to a scaling operation on the data domain, by dividing the expressions with a quantity proportional to $\sum_x G(x)$. The proportionality factor can be adjusted in order to determine the crispness of the decision.

3.4 Brownian Noise & Teager Energy

Another assumption made throughout the previous analysis was that the error between the model prediction and the observed signal can be modelled as white Gaussian noise (WGN). For correlated noise the parameter estimation formulae involve diagonalizing the noise covariance matrix and become fairly complicated [18, 7]; for the special case of Brownian motion, however, one can work on the derivative of the observations which can be modelled in terms of the derivatives of the basis elements plus WGN. Specifically, for model I_T which includes a single sinusoid of frequency ω_0 , this amounts to projecting the differentiated signal O' on the differentiated basis elements, $\mathcal{B}'_{T,o} = \omega_0 \mathcal{B}_{T,e}(x)$ and $\mathcal{B}'_{T,e} = -\omega_0 \mathcal{B}_{T,o}$; normalization leads to the omission of the multiplicative ω_0 factors, so the estimated amplitude A' for the differentiated signal equals

$$\begin{aligned} A' &= \sqrt{(O' * G\mathcal{B}_{T,e})^2 + (O' * G\mathcal{B}_{T,o})^2} \\ &= \sqrt{(O * G\mathcal{B}'_{T,e})^2 + (O * G\mathcal{B}'_{T,o})^2} \\ &= \omega_0 \sqrt{(O * G\mathcal{B}_{T,o})^2 + (O * G\mathcal{B}_{T,e})^2} = \omega_0 A. \end{aligned}$$

Iterating the previous analysis, the data likelihood under the textured signal hypothesis can be expressed in terms of the quantity $(A')^2 = \omega_0^2 A^2$, where A is the amplitude estimate for the non-differentiated signal. The quantity $\omega_0^2 A^2$ equals the *Teager Energy* of the sinusoidal, which we have found empirically in [9] to give better results for DCA [6] in the channel selection procedure than the typically used amplitude-based criterion. The above approach justifies its use in a probabilistic setting, and has allowed its integration with the pure amplitude-based selection criterion; details will be given in a larger version of this paper.




Figure 1. Discrimination of textured areas from edges. Top row: (a) Input image, (b) texture and (c) edge model amplitude estimates, respectively and model-based probabilities of (d) smooth, (e) textured and (f) edge regions respectively. Bottom row: (g)/(j) Input images and posterior probabilities of (h)/(k) texture and (i)/(l) edge models respectively.

4 Experimental Results

For all the results presented in this section, we have used a Gabor filterbank consisting of 40 isotropical filters at 10 orientations and 4 scales logarithmically placed on the frequency domain; the spread of the Gaussian is set equal to the half of the modulating sinusoid's period. For each Gabor filter a corresponding edge detection QFP has been used, detecting a phase congruent signal of the form of Eqn. (1). The Fourier series coefficients in Eqn. (1) have been set equal to those of a periodic square wave, whose period is four times that of the sinusoid used for the texture model. The confidence function $G(x)$ for the edge-detection QFP and the DC model is taken to be equal to the Gaussian of the Gabor filter, thereby rendering the energy measurements directly comparable.

The proposed approach has been applied on a variety of natural images like those of Fig. 1 giving plausible results. In these images the probability of an edge is typically higher along the borders of objects, while at textured regions the probability of an edge is lower than what would be indicated by a direct application of an edge detection filter. We also observe that along the borders of objects there is a decrease in the probability of texture, since the edge model explains away the intensity variation.

4.1 Edge Detection

Using our model we can provide an edge detection algorithm with a confidence measure in any measured edge characteristic. In order to verify this we performed a small-scale test on the benchmark of [14], where ground-truth data are available. The output of the QFP used for edge detection was multiplied with the posterior probability of an edge, calculated according to our model; as can be seen in Fig. 2, this helps suppress edges at textured areas, while keeping most true edges intact. Quantitative results on the test set

(100 images) are also provided, showing a systematic improvement in detection performance.

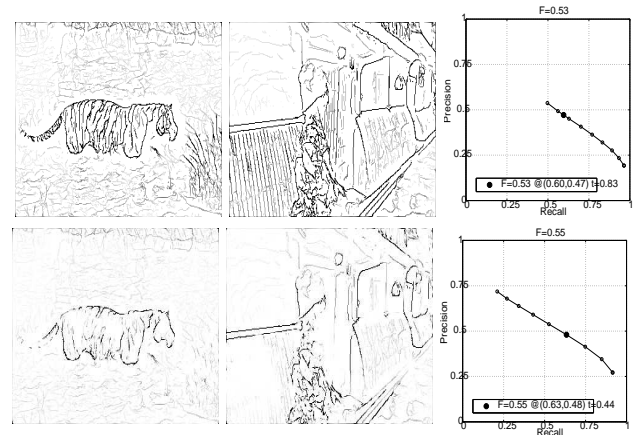


Figure 2. Edge detection results using oriented energy (top row) and after gating with the probability of edge (bottom row). In the last column comparative results on the Berkeley test set [14] are shown: False alarms are systematically eliminated, yielding higher precision values at the same recall rates.

4.2 Texture Segmentation

We can combine these ideas with the Dominant Components Analysis (DCA) [6] model for texture analysis which provides useful features for the segmentation of textured images [17, 10, 4]. This model picks at each image pixel the output of the strongest amongst a set of narrow-band filters, demodulates its output and uses the estimated AM-FM signals as texture features; we can view this as initially modeling the image in terms of Eqn. (2), and subsequently refining the estimated amplitude and frequency measurements. This model can capture a large part of the image

variation of textured images; due to lack of space, the interested reader is referred to [6, 9] for details on the feature extraction algorithms related to this model and [17, 10, 4] for results in unsupervised texture image segmentation. Even though this model gives informative features on textured areas, on smooth regions they are meaningless: for example the texture orientation signal used in [17, 10] is an informative texture cue but behaves erratically on smooth regions. Further, the amplitude estimates are high wherever edges appear, preventing the segmentation process from accurately detecting the object borders.

The posterior probabilities assigned to each hypothesis can be used to improve feature-based segmentation algorithms, by taking into account the confidence associated with each feature. In our case we have modified the Region Competition [20] curve evolution algorithm we used in [9] as follows: in the original algorithm, the boundary of region i evolves according to equation

$$\frac{\partial C_i}{\partial t} = \mathcal{N}_i \log \frac{P_i(F)}{P_j(F)} - \mathcal{N}_i \kappa_i, \quad (11)$$

where F is the set of features used to drive the segmentation process, κ_i is the front's curvature, \mathcal{N}_i is the front's outward normal vector, P_i is the probability distribution for the features inside region i and P_j is that of the competing neighboring region, j . This evolution equation drives each region i to the data that it can best model, according to its distribution, P_i . Using the model posterior probabilities, we modify this evolution law by taking into account the confidence in the features of each hypothesis i as follows [4]:

$$\frac{\partial C_i}{\partial t} = \mathcal{N}_i \sum_{h \in H} w_h \log \frac{P_{h,i}(F_h)}{P_{h,j}(F_h)} - \mathcal{N}_i \kappa_i, \quad (12)$$

where w_h is weight associated with hypothesis i . According to this evolution law, each region entertains simultaneously two distributions, one for the textured and one for the non-textured (e.g. intensity) features. Whenever the texture hypothesis prevails, its features can be used to drive the region competition process and vice versa. This modified evolution equation is extensively presented and justified in [4], where further results and a quantitative evaluation of its merits can be found. In Fig. 3 we show indicative results comparing the original and the modified evolution equations; in all cases, the fused features give better segmentations, with the segment borders accurately locating the object borders, where the texture features are suppressed.

5 Conclusion

In this work we have presented how simple generative models can discriminate between broad types of features, texture and edges. Our experimental results demonstrate its applicability to a wide variety of natural images, where




Figure 3. Segmentation results using the original (left) and the modified (right) evolution equation.

it can be used to offer a confidence estimate on the features used for subsequent tasks, like texture segmentation and edge detection.

References

- [1] BOUKERROULI, D., NOBLE, J., AND BRADY, M. On the Choice of Band-Pass Quadrature Filters. *Jrnl. of Math. Im. & Vision* 21 (2004), 53–80.
- [2] BOVIK, A., CLARK, M., AND GEISLER, W. Multichannel Texture Analysis using Localized Spatial Filters. *IEEE Trans. PAMI* 12, 1 (1990), 55–73.
- [3] BOVIK, A. C., GOPAL, N., EMMOTH, T., AND RESTREPO, A. Localized Measurement of Emergent Image Frequencies by Gabor Wavelets. *IEEE Trans. Information Theory* 38 (1992), 691–712.
- [4] EVANGELOPOULOS, G., KOKKINOS, I., AND MARAGOS, P. Advances in Variational Image Segmentation using AM-FM models: Regularized Demodulation and Probabilistic Cue Integration. In *Proc. 3rd. Intl. Workshop on Variational and Level Set Methods (VLSM)* (2005).
- [5] GUO, C.-E., ZHU, S.-C., AND WU, Y. N. A Mathematical Theory of Primal Sketch and Sketchability. In *Proc. ICCV* (2003).
- [6] HAVLICEK, J., AND BOVIK, A. Image Modulation Models. In *Handbook of Image and Video Processing*, A. Bovik, Ed. 2000, pp. 305–316.
- [7] KAY, S. M. *Fundamentals of Statistical Signal Processing: Estimation & Detection Theory*, vol. I&II. Prentice Hall, 1993.
- [8] KNUTSSON, H., AND WESTIN, C. F. Normalized Convolution and Differential Convolution: Methods for the Interpolation and Filtering of Incomplete and Uncertain Data. In *IEEE Conf. CVPR* (1993).
- [9] KOKKINOS, I., EVANGELOPOULOS, G., AND MARAGOS, P. Advances in Texture Analysis: Energy Dominant Components and Multiple Hypothesis Testing. In *Int'l Conf. on Image Processing* (2004).
- [10] KOKKINOS, I., EVANGELOPOULOS, G., AND MARAGOS, P. Modulation-Feature based Textured Image Segmentation Using Curve Evolution. In *Int'l Conf. on Image Processing* (2004).
- [11] LINDBERG, T. Edge Detection and Ridge Detection with Automatic Scale Selection. *Int'l. J. of Comp. Vision* 30, 2 (1998).
- [12] MALLAT, D., AND ZHANG, Z. Matching Pursuit in a Time-Frequency Dictionary. *IEEE Trans. Signal Processing* 41 (1993), 3397–3415.
- [13] MARTIN, D., FOWLKES, C., AND MALIK, J. Learning to Detect Natural Image Boundaries Using Local Brightness, Color, and Texture Cues. *IEEE Trans. PAMI* 26, 5 (2004), 530–549.
- [14] MARTIN, D., FOWLKES, C., TAL, D., AND MALIK, J. A Database of Human Segmented Natural Images and its Application to Evaluating Segmentation Algorithms and Measuring Ecological Statistics. In *Proc. ICCV* (2001).
- [15] MORRONE, C., AND BURR, D. Feature Detection in Human Vision: a Phase-Dependent Energy Model. *Proc. R. Soc. London B* 235 (1988), 221–245.
- [16] PERONA, P., AND MALIK, J. Detecting and Localizing Edges Composed of Steps, Peaks and Roofs. In *Proc. ICCV* (1990).
- [17] TANGSUKSON, T., AND HAVLICEK, J. P. AM-FM Image Segmentation. In *Int'l Conf. on Image Processing* (2000).
- [18] TREES, H. V. *Detection Estimation and Modulation Theory*. Wiley, 1971.
- [19] WESTIN, C. F., NORDBERG, K., AND KNUTSSON, H. On the Equivalence of Normalized Convolution and Normalized Differential Convolution. In *IEEE Conf. CVPR* (1994).
- [20] ZHU, S., AND YUILLE, A. Region Competition: Unifying Snakes, Region Growing and Bayes/MDL for Multiband Image Segmentation. *IEEE Trans. PAMI* 18, 9 (1996), 884–900.