AM-FM IMAGE ANALYSIS TECHNIQUES

J.P. Havlicek, M.S. Pattichis, D.S. Harding, A.C. Christofides, and A.C. Bovik

Laboratory for Vision Systems
University of Texas, Austin, TX 78712-1084

ABSTRACT

We present an overview of AM-FM image analysis techniques. Many nonstationary textured images are multi-partite, and may be considered to comprise several locally narrowband components. We introduce a multi-component AM-FM model useful in describing such images, and show that the components may be demodulated on a spatially local basis using multiband filtering followed by a computationally efficient nonlinear algorithm. We outline practical approaches both for estimating the locally dominant image frequencies and for computing a full AM-FM image representation. Finally, we discuss how the estimated frequencies may be used to compute flow lines.

I. INTRODUCTION

AM-FM techniques model images with nonlinear functions of the form

$$t(\mathbf{x}) = a(\mathbf{x}) \exp\left[j\varphi(\mathbf{x})\right], \tag{1}$$

which are inherently capable of capturing nonstationary features and structure. Such techniques offer advantages for computing efficient representations, and facilitate analysis with respect to the nonstationarities. AM-FM modeling is most useful for treating images that may be decomposed into *locally coherent* components, each admitting modulating functions $a(\mathbf{x})$ and $\nabla\varphi(\mathbf{x})$ which are smooth in the sense of having bounded Sobolev norms.

A single component of the form (1) may be demodulated using the local nonlinear algorithm

$$a(\mathbf{x}) = |t(\mathbf{x})|, \tag{2}$$

$$\nabla\varphi(\mathbf{x}) = \mathrm{Re}\left[\frac{\nabla t(\mathbf{x})}{jt(\mathbf{x})}\right], \tag{3}$$

which is *exact* for *any* complex AM-FM component [1].

Although any image may be exactly represented by a single AM-FM component, often no such representation will admit smooth modulating functions. Therefore, it becomes desirable to employ multi-component models of the form

$$t(\mathbf{x}) = \sum_{k=1}^{K} a_k(\mathbf{x}) \exp[j\varphi_k(\mathbf{x})]. \tag{4}$$

Real-valued images may be analyzed against the models (1), (4) by using the 2D Hilbert transform to create a complex-valued extension. To estimate $a_k(\mathbf{x})$ and $\nabla\varphi_k(\mathbf{x})$, we analyze the image with a multiband filterbank to separate components on a pixel-wise basis. For a filter with impulse response $g_m(\mathbf{x})$ and Frequency response $G_m(\boldsymbol{\Omega}) = \mathcal{F}[g_m(\mathbf{x})]$, the filtered component

$$t_m(\mathbf{x}) = \int_{\mathbb{R}^n} t(\mathbf{x}-\mathbf{p}) g_m(\mathbf{p}) d\mathbf{p} \tag{5}$$

may be demodulated using the approximate algorithm

$$\nabla\varphi(\mathbf{x}) \approx \nabla\widehat{\varphi}(\mathbf{x}) = \mathrm{Re}\left[\frac{\nabla t_m(\mathbf{x})}{jt_m(\mathbf{x})}\right], \tag{6}$$

$$a(\mathbf{x}) \approx \widehat{a}(\mathbf{x}) = \left|\frac{t_m(\mathbf{x})}{G_m\left[\nabla\widehat{\varphi}(\mathbf{x})\right]}\right|. \tag{7}$$

The approximation errors are generally small if $g_m(\mathbf{x})$ is spatially localized and the modulating functions are locally coherent [1,2].

II. DOMINANT COMPONENT ANALYSIS

Many images arising from natural physical, chemical, biological, and erosive processes contain textured regions or quasi-repetitive structures that can be described by the *emergent* frequencies that dominate the local image spectrum. Examples include crystals, rock strata, a zebra's stripes, wind patterns in sand, and wood grains. Perspective distortion may also give rise to nonstationary, quasi-regular patterns. Often, it is desirable to characterize the orientation, roughness, or flow of such patterns by estimating the dominant local frequencies.

In this section we describe an analysis technique known as *dominant component analysis* which, for an image of the form (4), will estimate the quantities $\nabla\varphi(\mathbf{x})$ and $a(\mathbf{x})$ of the locally dominant component; the frequencies $\nabla\widehat{\varphi}(\mathbf{x})$ then estimate the emergent image frequencies. A block diagram of the approach is shown in Figure 1, where the image $t(\mathbf{x})$ is ana-

This research was supported in part by a grant from the Texas Advanced Research Projects Agency and by the Air Force Office of Scientific Research, Air Force Systems Command, USAF, under grant number F49620-93-1-0307.

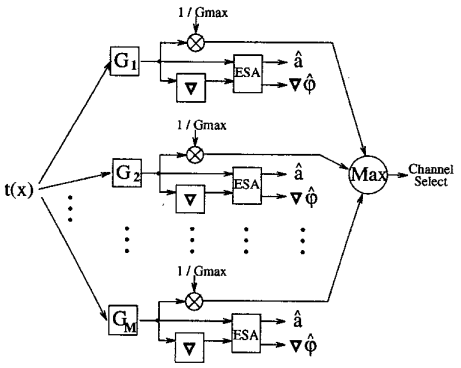

Figure 1: Block diagram of the dominant component technique. The demodulation algorithm (6), (7) is performed in the blocks marked ESA.

lyzed with an M-channel multiband bank of Gabor filters arranged in a wavelet-like tesselation [2]. The filtered demodulation algorithm (6), (7) is applied to the response of each channel at each pixel. Estimates $\nabla\widehat{\varphi}(\mathbf{x})$, $\widehat{a}(\mathbf{x})$ for the locally dominant component are taken from the channel maximizing the filter selection criterion

$$\Psi_m(\mathbf{x}) = \frac{|t_m(\mathbf{x})|}{\max_{\Omega} |G_m(\Omega)|}. \quad (8)$$

Figure 2 shows dominant component analysis of two Brodatz images. The image *raffia* is shown in Figure 2(a). The emergent frequency estimates are shown in the (b) part of the figure, where arrow length is proportional to the instantaneous period and the arrow heads point in the direction of $\nabla\widehat{\varphi}(\mathbf{x})$. Hence shorter arrows correspond to high frequencies and image features of small spatial extent, while longer arrows correspond to features of larger spatial extent. Although this image clearly has a multi-partite nature, the strips oriented with the main antidiagonal are particularly strong throughout the image, and this is captured in the emergent frequency estimates. The image *tree* is shown in Figure 2(c). The emergent frequency estimates are shown in the (d) part of the figure, where the lengths of the arrows are squared to accentuate the differences between the highest and lowest frequency estimates.

III. MULTI-COMPONENT AM-FM REPRESENTATION

In this section, we outline a practical approach for computing the multi-component AM-FM representation

$$\{\widetilde{a}_k(m,n), \nabla\widetilde{\varphi}_k(m,n)\}_{\substack{k\in[1,K] \\ (m,n)\in[0,M-1]\times[0,N-1]}} \quad (9)$$

of an M-column by N-row K-component discrete-domain image modeled as samples of (4). This re-

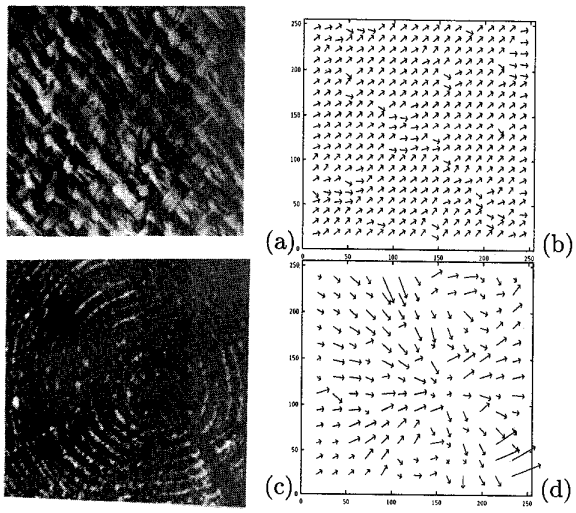

Figure 2: Dominant component analysis. (a) *raffia*. (b) Needle diagram showing $\nabla\varphi(\mathbf{x})$ for *raffia*. (c) *tree*. (d) Needle diagram showing $\nabla\varphi(\mathbf{x})$ for *tree*

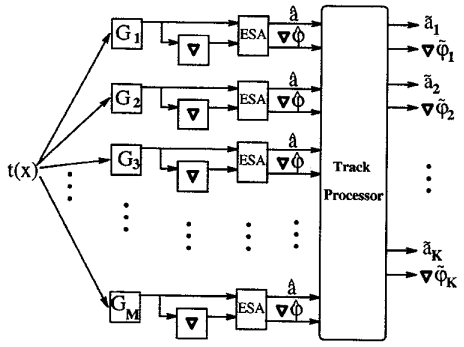

Figure 3: Block diagram of the approach for computing Multi-component AM-FM representations.

quires simultaneously estimating the modulating functions of *all* components. Figure 3 shows a block diagram of the computational paradigm. As before, components are isolated on a spatio-spectrally local basis using a multiband filterbank. Subsequent to filtering, the demodulation algorithm (6), (7) is applied to each channel response. The track processor in Figure 3 assimilates the estimates obtained from each channel into estimates for each component to obtain the representation (9). In Sections IIIA - IIIC. we describe a novel track processor based on Kalman filters derived from a statistical state-space component model.

A. Statistical Component Model

Assume an image of the form (4), and define $\varphi^x(\mathbf{x}) = \frac{\partial}{\partial x}\varphi(\mathbf{x})$. Introduce an artificial temporal causality

relationship between points in the sampled image domain by mapping them to a 1D lattice according to a path function $\mathcal{O} : \mathbf{x} \longmapsto k$, $k \in \mathbb{N}$. $\mathcal{O}$ maps the modulating function of a component according to, *e.g.* $a(\mathbf{x}) \stackrel{\mathcal{O}}{\longmapsto} a(k)$. Let ρ denote continuous arc length along $\mathcal{O}$, and use the notation

$$\varphi^{x\prime}(k) = \frac{\partial}{\partial \rho}\varphi^x(\rho)\bigg|_{\rho=k}, \tag{10}$$

to indicate the restriction to the discrete 1D lattice of the derivative of a modulating function taken with respect to ρ. Then each modulating function may be expanded in a first-order Taylor series about a lattice point k, *e.g.*

$$\begin{aligned}\varphi^x(k+1) &= \varphi^x(k) + \varphi^{x\prime}(k) \\ &\quad + \int_k^{k+1} (k+1-\rho)\frac{\partial^2}{\partial \rho^2}\varphi^x(\rho)d\rho. \end{aligned}\tag{11}$$

Likewise, the first-order derivatives of the modulating functions may be expanded in zeroth-order Taylor series:

$$\varphi^{x\prime}(k+1) = \varphi^{x\prime}(k) + \int_k^{k+1} \frac{\partial^2}{\partial \rho^2}\varphi^x(\rho)d\rho. \tag{12}$$

We consider that $a(\mathbf{x})$ and $\varphi(\mathbf{x})$ are independent, homogeneous m.s. differentiable random fields, and that $\varphi(\mathbf{x})$ is *quadrant symmetric* [3]. Then we model the Taylor series integrals associated with (11) by three noise processes $u_a(k)$, $u_{\varphi_x}(k)$, and $u_{\varphi_y}(k)$. The Taylor series integrals associated with (12) are modeled by noise processes $\nu_a(k)$, $\nu_{\varphi_x}(k)$, and $\nu_{\varphi_y}(k)$. Collectively, we refer to these six noise processes as the *modulation accelerations*, or *MA's*, since they involve local averages of the second derivatives of the modulating functions. Then, the series (11), (12) for all three modulating functions may be written together in a canonical state-variable form to obtain the statistical state-space component model

$$\begin{bmatrix} a(k+1) \\ a'(k+1) \\ \varphi^x(k+1) \\ \varphi^{x\prime}(k+1) \\ \varphi^y(k+1) \\ \varphi^{y\prime}(k+1) \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} a(k) \\ a'(k) \\ \varphi^x(k) \\ \varphi^{x\prime}(k) \\ \varphi^y(k) \\ \varphi^{y\prime}(k) \end{bmatrix}$$
$$+ \left[u_a(k)\ \nu_a(k)\ u_{\varphi_x}(k)\ \nu_{\varphi_x}(k)\ u_{\varphi_y}(k)\ \nu_{\varphi_y}(k) \right]^T. \tag{13}$$

We model estimation errors in the algorithm (6), (7) by noise processes $n_a(k)$, $n_{\varphi_x}(k)$, and $n_{\varphi_y}(k)$, called the *measurement noises*, or *MN's*. This results in the observation equation

$$\begin{bmatrix} \widehat{a}(k) \\ \widehat{\varphi^x}(k) \\ \widehat{\varphi^y}(k) \end{bmatrix} = \begin{bmatrix} a(k) \\ \varphi^x(k) \\ \varphi^y(k) \end{bmatrix} + \begin{bmatrix} n_a(k) \\ n_{\varphi_x}(k) \\ n_{\varphi_y}(k) \end{bmatrix}. \tag{14}$$

We assume that the covariance structures of the second derivatives of the modulating functions are highly spatially localized, so that they behave impulsively in both the vertical and horizontal directions when viewed at the scale of the spatial sampling lattice. Then, it is straightforward to compute the following covariances between the MA's:

$$E[u_a(k)u_a(j)] = \frac{1}{3}\sigma^2_{a''}\delta(k-j), \tag{15}$$
$$E[\nu_a(k)\nu_a(j)] = \sigma^2_{a''}\delta(k-j), \tag{16}$$
$$E[u_a(k)\nu_a(j)] = \frac{1}{2}\sigma^2_{a''}\delta(k-j). \tag{17}$$

Expressions for $E[u_{\varphi_x}(k)u_{\varphi_x}(j)]$, $E[u_{\varphi_y}(k)u_{\varphi_y}(j)]$, $E[\nu_{\varphi_x}(k)\nu_{\varphi_x}(j)]$, $E[\nu_{\varphi_y}(k)\nu_{\varphi_y}(j)]$, $E[u_{\varphi_x}(k)\nu_{\varphi_x}(j)]$, and $E[u_{\varphi_y}(k)\nu_{\varphi_y}(j)]$ are analogous in form. The remaining twelve MA covariance functions are all identically zero.

Analytical characterization of the MN moments and of the joint moments between the MN's and MA's is difficult. However, in practice we have found that when dealing with locally coherent images, the MN moments are normally quite small with relation to the magnitudes of the modulating functions themselves. Therefore, within the scope of this paper we shall assume that the MN's are jointly uncorrelated and also uncorrelated with the MA's.

B. Component Track Filter

In the previous section, we modeled the estimated modulating functions of a component as noisy observations of an affine function of the state vector of a finite-dimensional linear system driven by uncorrelated noise. Hence, the MMSE optimal track processor involves Kalman filters. Due to the block-diagonal structure of the state transition matrix in (13), the system modes corresponding to the amplitude modulation and the two components of the frequency modulation may be decoupled, resulting in three independent second-order systems. We use independent Kalman filters to track each one.

The explicit recursive formulation for the optimal amplitude estimates $\widetilde{a}(k|k)$ is

$$\begin{aligned}\widetilde{a}(k|k) &= \widetilde{a}(k|k-1) \\ &\quad + \alpha_a(k)[\widehat{a}(k) - \widetilde{a}(k|k-1)], \end{aligned}\tag{18}$$
$$\begin{aligned}\widetilde{a}'(k+1|k) &= \widetilde{a}'(k|k-1) \\ &\quad + \beta_a(k)[\widehat{a}(k) - \widetilde{a}(k|k-1)], \end{aligned}\tag{19}$$
$$\widetilde{a}(k+1|k) = \widetilde{a}(k|k) + \widetilde{a}'(k+1|k). \tag{20}$$

The optimal frequency estimators are analogous in form, and involve gain sequences $\alpha_{\varphi_x}(k)$, $\beta_{\varphi_x}(k)$, $\alpha_{\varphi_y}(k)$, and $\beta_{\varphi_y}(k)$. Formulations for $\alpha_{\square}(k)$, $\beta_{\square}(k)$ follow from recursive expressions for the state vector

error covariance matrices associated with the decoupled systems. This track processor is unbiased and optimal in the MMSE sense.

C. New Track Starts

At each point i in the path $\mathcal{O}$, we start new tracks using the observations $\widehat{a}(i)$ and $\nabla\widehat{\varphi}(i)$ obtained from channels which maximize the quantity

$$Q(n) = \frac{\widehat{a}_n(i)\,|G_n[\nabla\widehat{\varphi}_n(i)]|}{\max_{\mathbf{\Omega}}|G_n(\mathbf{\Omega})|}, \tag{21}$$

where $\widehat{a}_n(i)$ and $\nabla\widehat{\varphi}_n(i)$ are the amplitude and frequency estimates obtained from filterbank channel n, provided that these observations do not associate with an already existing track. Note that the quantity $|G_n(\mathbf{\Omega})| \,/\, \max_{\mathbf{\Omega}}|G_n(\mathbf{\Omega})|$ lies between zero and one, and increases as $\mathbf{\Omega}$ moves closer to the center frequency of the channel. Hence, for a given component, tracking will be initialized using the channel whose center frequency is closest to the instantaneous frequency of the component, affording improved immunity against out-of-band information through an enhanced SNR.

The amplitude track is initialized at pixel i using channel n by taking

$$\widetilde{a}(0|0) = \widehat{a}_n(i) \tag{22}$$

and

$$\widetilde{a'}(0|0) = \widehat{a}_n(i) - \widehat{a}_n(i-1). \tag{23}$$

Similar procedures are used to initialize the frequency tracks. Typically, for some small number L, new tracks are started at each pixel using the amplitude and frequency observations from the L channels having the largest values of $Q(n)$ from among those channels producing observations that do not associate with an already existing track. In an image region where K components are present, this scheme guarantees that *all* components will be tracked by the $\lceil K/L \rceil^{\text{th}}$ pixel of the region (under the assumption that at most one component dominates the local frequency spectrum of each channel at each pixel).

D. Multi-Component Example

In Figure 4 we compute the representation (9) of the synthetically generated, two-component nonstationary image shown in Figure 4(a). True values for components $C1$ and $C2$ are shown in Figures 4(b) and (c), respectively. The presence of two components was correctly identified by the track processor. The amplitude estimates for $C1$ are shown in Figure 4(d), while the horizontal and vertical frequency estimates are given in the (e) and (f) parts of the figure, respectively. For $C2$, the amplitude, horizontal frequency, and vertical frequency estimates are given in the (g), (h), and (i) parts of the figure. Collectively, Figures 4(d) – (i) constitute the multi-component AM-FM representation of the image. Note how smooth the AM-FM representation is, despite the fact that there are rapid variations in the image. The estimated quantities are in near perfect agreement with the true values, except for a small region of oscillatory behavior in $C1$ where the vertical frequency approaches DC.

IV. A STOCHASTIC FLUID MODEL

In this section, we use the emergent frequency vectors to develop a stochastic fluid model which provides an interpretation of frequency modulation. Here, we make an analogy between emergent frequency vectors and particle velocities [4]. In [4], this analogy lead to a deterministic fluid model that was used to analyze grain variations in a woodgrain texture. In the stochastic fluid model considered here, we show that the frequency modulation process can be described in terms of two independent free Brownian motion models [5]. We show that particle trajectories, or *flow lines*, are samples of Brownian motions. Frequency modulation is analyzed in terms of the parameters of Brownian motions.

We use the notation $\mathbf{O}$ to denote the emergent frequency vector $\nabla\varphi(\mathbf{x})$. We consider $\mathbf{O}$ to be a velocity field, and $\mathbf{x}(t)$ to be a particle trajectory of $\mathbf{O}$ [4]. For a continuously differentiable velocity field, we have by the chain rule that

$$\begin{bmatrix} dO_1/dt \\ dO_2/dt \end{bmatrix} = \begin{bmatrix} \partial O_1/\partial x & \partial O_1/\partial y \\ \partial O_2/\partial x & \partial O_2/\partial y \end{bmatrix} \begin{bmatrix} dx/dt \\ dy/dt \end{bmatrix}. \tag{24}$$

Since $\mathbf{O}$ is also the tangent to the trajectory $\mathbf{x}$, we have that $\mathbf{O}'(\mathbf{x}(t)) = \mathbf{F}(\mathbf{x})\,\mathbf{O}(\mathbf{x}(t))$, where $\mathbf{F}$ is the emergent frequency gradient tensor [4]. Hence, (24) describes the deterministic fluid model that was considered in [4]. We define a stochastic fluid model by incorporating a normal white noise process $\mathbf{w}(\mathbf{x}(t))$ to account for random disturbances along the particle trajectory [5]:

$$\mathbf{O}'(\mathbf{x}(t)) = \mathbf{F}(\mathbf{x}(t))\,\mathbf{O}(\mathbf{x}(t)) + \mathbf{w}(\mathbf{x}(t)) \tag{25}$$

In what follows, for slowly-varying $\mathbf{F}$, we show that (25) can be modeled by two independent free Brownian motions (as in [5]).

Next, we address certain estimation issues as they relate to (25). The emergent frequency vector estimates must be extended to four quadrants, as was done in [4]. We set $\mathbf{O} = \pm\nabla\varphi(\mathbf{x})$, choosing the sign to satisfy $\mathbf{O}\cdot\mathbf{O}' < 0$. This ensures that the emergent frequency vector magnitude will decrease along flow lines. In addition, we note that for a true gradient field, the emergent frequency gradient tensor must be

symmetric. Hence, we take $\frac{1}{2}(\partial^2\widehat{\varphi}/\partial x\partial y + \partial^2\widehat{\varphi}/\partial y\partial x)$ as the estimate for both $\partial O_1/\partial x$ and $\partial O_2/\partial y$.

Since the frequency gradient tensor is symmetric, $\mathbf{F}$ may always be expressed in terms of its eigenvalue-eigenvector decomposition, $\mathbf{F} = \mathbf{E}^T\mathbf{\Lambda}\mathbf{E}$. Expressing the emergent frequency vector and random disturbance in terms of the eigenvector coordinate system, we have $\mathbf{z} = \mathbf{EO}$ and $\mathbf{w}_n = \mathbf{Ew}$. More importantly, we make the approximation $\mathbf{z}' \approx \mathbf{EO}'$, which is valid for $||d\mathbf{E}/dt||$ small. It can be shown that, along flow lines, the approximation holds well when the unit eigenvectors $\mathbf{e}_i$ rotate slowly [4]. In view of this, (25) may be rewritten as

$$z_i' = \lambda_i z_i + (\mathbf{w}_n)_i \qquad i = 1, 2, \tag{26}$$

where λ_i are the eigenvalues of $\mathbf{F}$.

We are now ready to provide physical interpretations for all the components of the model. First, we note that (26) describes two independent free Brownian motions [5]. Each flow line corresponds to a sample from an ensemble of Brownian motions. Here, z_i denotes the particle velocity (emergent frequency vector component) of a small particle moving in the direction of the eigenvector $\mathbf{e}_i$. During the motion, each particle velocity component z_i is reduced by friction forces. In our model, the friction coefficients are given by the negative eigenvalues λ_i. In addition, the motion is affected by the random collision forces $(\mathbf{w}_n)_i$. Furthermore, $(\mathbf{w}_n)_i$ is a normal white noise, since the collision forces along the motion are uncorrelated and the forces are normally distributed (with zero mean). The variance is directly proportional to the friction coefficient and the *temperature* of the medium.

A. The Discrete Model

To solve (26) for $\mathbf{z}$, we assume that $\mathbf{F}$ is piecewise constant. This leads to the Kalman filter system

$$\mathbf{z}_{k+1} = \mathbf{A}_k\mathbf{z}_k + \mathbf{B}_k\,(w_n)_k \tag{27}$$

$$\mathbf{y}_k = \mathbf{z}_k + \mathbf{v}_k, \tag{28}$$

where $\mathbf{y}_k$ denotes a measurement of $\mathbf{z}_k$ corrupted by additive noise $\mathbf{v}_k$. It may be shown that

$$\mathbf{A}_k = \exp\{\,\mathbf{\Lambda}_k\,\}, \tag{29}$$

$$\mathbf{B}_k = \begin{bmatrix} b_{1,1} & 0 \\ 0 & b_{2,2} \end{bmatrix}\mathbf{E}_k, \tag{30}$$

where

$$b_{i,i} = \left(-\frac{1}{\lambda_i} + \frac{\exp\{\,\lambda_i\,\}}{\lambda_i}\right). \tag{31}$$

In Figure 5, we show an example of the discrete flow line model. The flow lines were generated using (27) - (28). Then, a Kalman filter was used to track noisy measurements of the emergent frequency vectors. The flow lines converge centrally, where the components of the emergent frequency vectors are nearly zero. This example demonstrates that the model is robust against noise in the frequency estimates.

V. CONCLUSION

In this paper, we demonstrated how complicated, nonstationary images may be modeled using multi-component AM-FM functions. We discussed multi-band filtering techniques for isolating the components on a spatio-spectrally localized basis, and presented an algorithm for demodulating the filtered components. Our approach bears similarities to processing known to occur in biological vision systems, in that the *information content* of an image is characterized by smoothly varying modulations occurring in a small number of frequency and orientation selective channels.

We gave practical computational techniques for estimating the emergent frequencies that dominate the local image spectrum and for computing multi-component AM-FM image representations. There are two main advantages in doing this. First, complicated nonstationary images may be represented by a small number of smoothly varying modulating functions which can be transformed to a spatially efficient form. Second, AM-FM representations naturally facilitate analysis *in terms of* the nonstationarities. Finally, we discussed how the estimated frequency modulations may be used to compute flow lines which characterize the orientation and flow of textured patterns.

VI. REFERENCES

[1] J. P. Havlicek, D. S. Harding, and A. C. Bovik, "Reconstruction from the multi-component AM-FM image representation", in *Proc. IEEE Int'l. Conf. Image Proc.*, Washington, DC, October 22-25 1995, pp. II280 - II283.

[2] A. C. Bovik, N. Gopal, T. Emmoth, and A. Restrepo, "Localized measurement of emergent image frequencies by Gabor wavelets", *IEEE. Trans. Info. Theory*, vol. IT-38, no. 2, pp. 691–712, March 1992.

[3] E. Vanmarcke, *Random Fields, Analysis and Synthesis*, MIT Press, Cambridge, MA, 1983.

[4] Marios S. Pattichis and Alan C. Bovik, "Multi-dimensional frequency modulation in texture images", in *International Conference on Digital Signal Processing*, Limassol, Cyprus, june 1995.

[5] Athanasios Papoulis, *Probability, Random Variables, and Stochastic Processes*, McGraw-Hill, New York, 3rd edition, 1991.

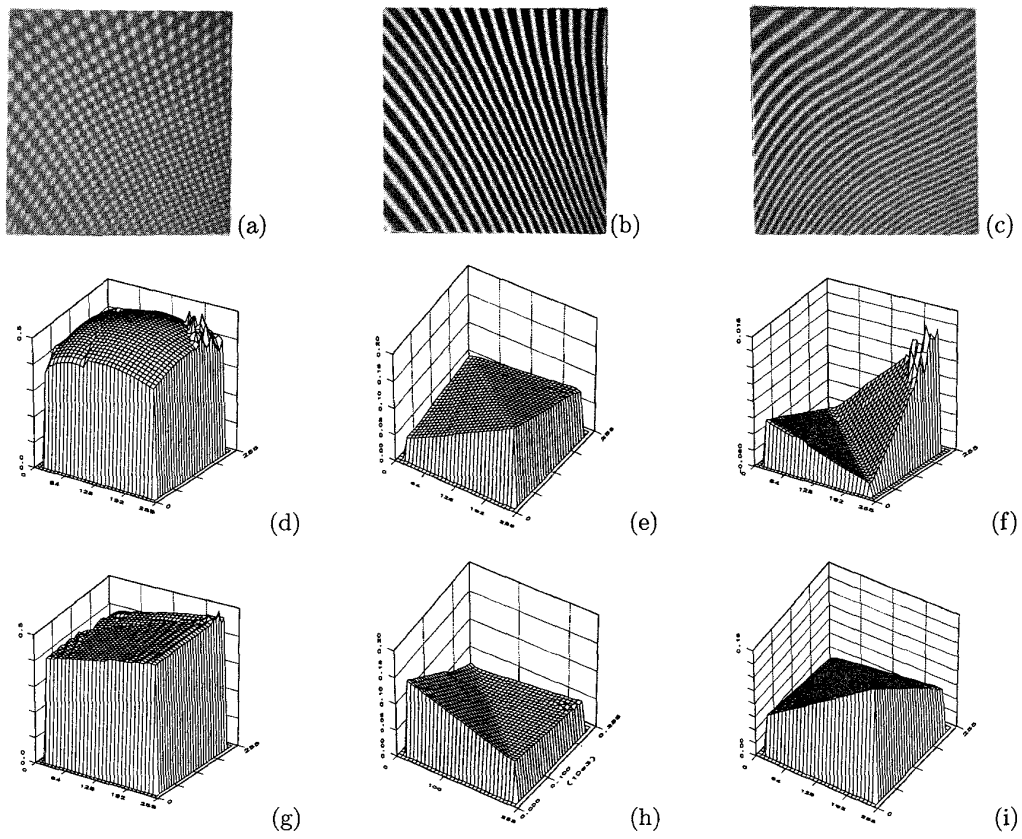

Figure 4: Multi-component AM-FM representation. (a) Nonstationary two-component synthetic image. (b) Component $C1$. (c) Component $C2$. (d) $\widetilde{a}_1(k)$. (e) $\widetilde{\varphi_1^x}(k)$. (f) $\widetilde{\varphi_1^y}(k)$. (g) $\widetilde{a}_2(k)$. (h) $\widetilde{\varphi_2^x}(k)$. (i) $\widetilde{\varphi_2^y}(k)$.

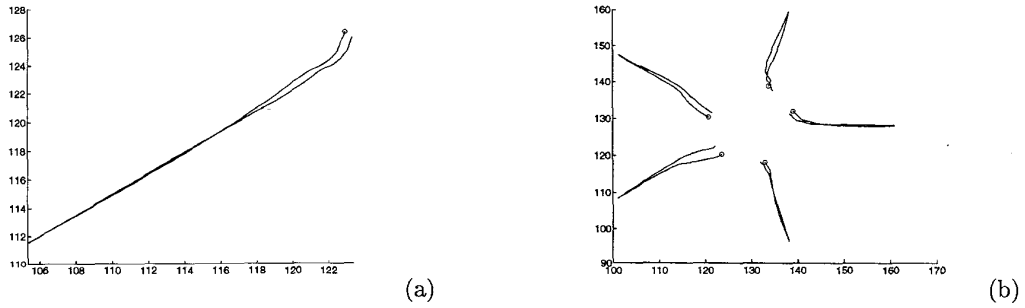

Figure 5: The stochastic fluid model (27), (28). (a) A typical flow line tracking example. (b) A collection of flow lines over the image plane. Both the estimated and the actual flow lines are shown. In both figures, estimated flow lines are terminated by the symbol 'o'. For the simulation, fifty values uniformly sampled between -0.1 and -0.01 where used for the eignevalues λ_i. For the Brownian motion, each component of $\mathbf{w}_n$ had a variance of $\sigma_{w_n} = \pi/100$, and for the observation process each component of $\mathbf{v}$ had a variance of $\sigma_v = \pi/10$. The estimated state-transition matrix $\widehat{\mathbf{A}}_k$ for the Kalman filter tracker was modeled by $\mathbf{A}_k$ plus uncorrelated noise. Each component of $\widehat{\mathbf{A}}_k$ was normally distributed about the true value with a variance of $\sigma_{\mathbf{A}} = 0.05$. For each noise process, the Kalman tracker utelized estimated variances which where five times the true values.