

FM FILTERS FOR MODULATION DOMAIN IMAGE PROCESSING

Chuong T. Nguyen, Patrick A. Campbell, and Joseph P. Havlicek

School of Electrical and Computer Engineering, University of Oklahoma, Norman, OK, USA

ABSTRACT

For the first time, we demonstrate modulation domain image filters that achieve perceptually motivated image processing goals by directly manipulating the FM functions in a multi-component AM-FM image model. The action of previous modulation domain filters has been limited to modification of the AM functions based on the values of the AM and FM functions. This is because reconstruction of the modified phase from the filtered frequency modulation vectors was an unsolved problem. Here, we present two new algorithms capable of reconstructing the phase from the processed frequencies, one based on a least squares solution of the discrete Poisson equation with Neumann boundary condition and one based on cubic tensor product spline integration. New modulation domain FM filters are designed to modify both the orientations and magnitudes of the visually important emergent image frequency vectors. In our most dramatic example, we demonstrate an FM filter that autonomously changes the stripes on the pants in the well known Barbara image from vertical to horizontal.

Index Terms— AM-FM image models, AM-FM image filters, modulation domain signal processing, multicomponent models

1. INTRODUCTION

Multicomponent AM-FM image models provide a perceptually motivated image representation from which nonstationary amplitude and frequency information can be readily interpreted. The model represents an image $t: \mathbb{R}^2 \rightarrow \mathbb{R}$ as a sum of K nonstationary AM-FM components [1–4]

$$t(\mathbf{x}) = \sum_{k=1}^K t_k(\mathbf{x}) = \sum_{k=1}^K a_k(\mathbf{x}) \cos[\varphi_k(\mathbf{x})], \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^2$, $a_k(\mathbf{x}) \geq 0$, and $\varphi_k(\mathbf{x}): \mathbb{R}^2 \rightarrow \mathbb{R}$. Given the image $t(\mathbf{x})$, a computed AM-FM model consists of estimates of the K AM functions $a_k(\mathbf{x})$ which provide a local characterization of the local contrast and the K FM functions $\nabla\varphi_k(\mathbf{x})$ which provide a dense characterization of the local texture orientations and pattern spacings.

These models have been used to great effect in a variety of computer vision and analysis applications, including texture segmentation and classification, 3-D shape from texture, edge detection and image enhancement, texture-based stereopsis, fingerprint classification, content-based retrieval, and regeneration of occluded and damaged textures [5], as well as for human speech analysis [6, 7] and recently in infrared [8] and visible [9] target tracking. While modulation domain image models have proven useful in *analysis* applications, they have found relatively limited use in applications requiring

This work was supported in part by the U.S. Army Research Laboratory and the U.S. Army Research Office under grant W911NF-08-1-0293.

image *synthesis*, primarily due to the difficulties involved in reconstructing an image from its modulation domain representation.

The recent development of perfect reconstruction AM-FM image models [10] has opened the door to the possibility of performing perceptually based image filtering directly in the modulation domain. The general notion is that after the image has been modeled according to (1) and the AM and FM functions have been computed, then signal processing can be applied to modify the AM and FM functions to achieve a desired signal processing goal. The filtered image is then reconstructed from the processed modulations. Extremely promising preliminary examples were given in [11, 12], where we demonstrated dramatic results not readily achievable via LTI filtering techniques.

However, the modulation domain image filters demonstrated in [11, 12] were restricted to “AM-only” processing: while both the AM and the FM functions were generally involved in determining the action of the filter, only the AM functions were actually modified (*e.g.*, frequency selective attenuation of the amplitude). Modulation domain filters that actually modify the FM functions have not been considered previously. The reason is that the reconstruction of the component phases $\varphi_k(\mathbf{x})$ in (1) from the computed FM functions $\nabla\varphi_k(\mathbf{x})$ is extremely sensitive to noise and to errors and perturbations in both the frequency vectors and in the boundary conditions (integration constants) used to reconstruct the phase from the frequencies [13].

In this paper, we introduce two new practical techniques for reconstructing the processed image after modulation domain filtering has been applied to modify the FM functions. For the first time, we design sophisticated filters that directly modify the frequency modulations of the image to achieve dramatic, perceptually motivated signal processing goals. We briefly discuss the invertible AM-FM image transform in Section 2. In Section 3, we develop the new phase reconstruction techniques and use them to define powerful FM filtering operations. Conclusions are reserved for Section 4.

2. PERFECT RECONSTRUCTION AM-FM MODEL

We decompose the image $t(\mathbf{x})$ into a sum of components $t_k(\mathbf{x})$ as indicated in (1) using an adaptation of the steerable pyramid [14, 15]. As described in [12], we perform a standard steerable pyramid decomposition to four levels with eight orientations at each level. We further decompose the high pass residual from the first level into eight orientations for a total of $K = 40$ components $t_k(\mathbf{x})$ in (1). The steerable pyramid is attractive because it admits smoothness and joint localization properties that are similar to those of Gabor filterbanks in terms of producing locally coherent components that correspond well to human visual perception. In addition, the steerable pyramid offers two distinct advantages: (1) whereas the AM and FM functions must be estimated from the Gabor filter responses which are not equal to the components t_k in general, the steerable pyramid provides direct access to the components themselves which may then

be demodulated using exact continuous domain algorithms that do not suffer from approximation errors, and (2) because the pyramid channels are self inverting, interband aliasing is minimized in the synthesis filterbank, thereby providing near perfect reconstruction.

After decomposition, we construct a complex extension $z_k(\mathbf{x}) = t_k(\mathbf{x}) + j\mathcal{H}[t_k(\mathbf{x})]$ for each component, where $\mathcal{H}[\cdot]$ is the partial Hilbert transform. We fit the complex components z_k with splines and differentiate analytically [10, 13] to implement the *exact* demodulation algorithm [3]

$$\nabla \hat{\varphi}_k(\mathbf{x}) = \operatorname{Re} \left[\frac{\nabla z_k(\mathbf{x})}{j z_k(\mathbf{x})} \right] \quad (2)$$

$$a_k(\mathbf{x}) = |z_k(\mathbf{x})|. \quad (3)$$

Unfortunately, the frequency field $\nabla \hat{\varphi}_k$ in (2) generally fails to lie in the space of cubic tensor product splines, which is not closed under division. Therefore, while the frequencies (2) are mathematically valid, spline integration filters cannot be used to reconstruct a consistent phase function from them. As explained in [16], we use the frequencies (2) to guide a novel unwrapping algorithm [13] that delivers a cubic spline phase function φ_k with gradient $\nabla \varphi_k$ that is closest to $\nabla \hat{\varphi}_k$ in the least squares sense, where $\nabla \varphi_k$ is obtained by analytical differentiation of φ_k [10, 16].

3. PHASE RECONSTRUCTION FOR FM FILTERING

After the application of FM signal processing, the modified frequency fields must be integrated to reconstruct the phase function of each processed image component. Based on the spline framework introduced in [17, 18], an error-free algorithm for reconstructing the phase φ_k from the FM function $\nabla \varphi_k$ was given in [13]. However, in order to achieve perfect reconstruction, this algorithm requires *unmodified* initial conditions on the phase. Moreover, it reconstructs the phase on the first two columns by integrating the vertical frequencies and then performs phase reconstruction on the rest of the image by integrating the horizontal frequencies across each row. Thus, vertical frequency modifications off the first two image columns are not reflected in the reconstructed phase, whereas small perturbations of the phase initial conditions and frequencies on the first two columns tend to cause severe instability in the phase reconstruction.

In this section, we introduce two new phase reconstruction algorithms that effectively overcome these problems and allow for consistent phase reconstruction after FM signal processing has been performed. Let $\mathbf{p} = [m \ n]^T \in \mathbb{Z}^2$ and let $t(\mathbf{p})$ be the original $M \times N$ digital image containing the samples of $t(\mathbf{x})$ in (1). Denote the set of pixel sites in the digital image $t(\mathbf{p})$ by

$$\mathcal{O} = \{\mathbf{p} : 0 \leq m \leq M - 1 \text{ and } 0 \leq n \leq N - 1\} \subset \mathbb{Z}^2. \quad (4)$$

We will write $\nabla \varphi_k(\mathbf{p})$ to indicate the restriction of the FM function $\nabla \varphi_k(\mathbf{x})$ of component $t_k(\mathbf{x})$ in (1) and its complex extension $z_k(\mathbf{x})$ to the discrete pixel lattice \mathcal{O} . Let \mathcal{G} denote the FM signal processing operation, so that the samples of the filtered frequency field are given by $\nabla \tilde{\varphi}_k(\mathbf{p}) = \mathcal{G}[\nabla \varphi_k(\mathbf{p})]$. Then, for $k \in [1, K]$, the phase reconstruction problem is to obtain the modified phase functions $\tilde{\varphi}_k(\mathbf{p})$ given the filtered frequencies $\nabla \tilde{\varphi}_k(\mathbf{p})$.

3.1. Least Squares Algorithm

In this section, we introduce a least squares phase reconstruction algorithm based on the well-known discretized Poisson equation with Neumann boundary condition [19]. Let \tilde{U}_k and \tilde{V}_k be the components of the filtered frequency vector defined by $[\tilde{U}_k(\mathbf{p}) \ \tilde{V}_k(\mathbf{p})]^T =$

$\nabla \tilde{\varphi}_k(\mathbf{p})$. Define \tilde{U}_k and \tilde{V}_k to be the increments of the modified phase $\tilde{\varphi}_k$ so that $\tilde{U}_k(m, n) = \tilde{\varphi}_k(m, n + 1) - \tilde{\varphi}_k(m, n)$ and $\tilde{V}_k(m, n) = \tilde{\varphi}_k(m + 1, n) - \tilde{\varphi}_k(m, n)$. We seek the modified phase function $\tilde{\varphi}_k$ that minimizes

$$E_k = \sum_{\mathbf{p} \in \mathcal{O}} [\tilde{U}_k(\mathbf{p}) - \tilde{U}_k(\mathbf{p})]^2 + [\tilde{V}_k(\mathbf{p}) - \tilde{V}_k(\mathbf{p})]^2. \quad (5)$$

A frequency domain solution for (5) was proposed in [19] and is given by

$$\tilde{\varphi}_k(m, n) = \text{IDCT} \left[\frac{\text{DCT}[\sigma_k(m, n)]}{2(\cos(\pi m/M) + \cos(\pi n/N)) - 4} \right], \quad (6)$$

where $\sigma_k(m, n)$ is the discrete Laplacian $\sigma_k(m, n) = \tilde{U}_k(m, n) - \tilde{U}_k(m, n - 1) + \tilde{V}_k(m, n) - \tilde{V}_k(m - 1, n)$, and where ‘‘DCT’’ and ‘‘IDCT’’ denote the discrete cosine transform and its inverse, respectively. The solution in (6) is defined everywhere except at the first pixel $\mathbf{p} = [0 \ 0]^T$, where it is set to zero.

We applied this phase reconstruction algorithm to perform FM image filtering. The results are shown in Fig. 1. By changing the orientation angles of the FM functions, we can achieve a rotation of the image texture by an arbitrary angle in the image. Fig. 1(b) illustrates an example of FM filtering where we rotate the frequency orientations by $\pi/2$ rad. Geometric stretching effects can also be obtained by scaling the magnitudes of the FM vectors. To implement such a filter, we multiply the FM vectors by a positive constant α , e.g., $\alpha > 1$ produces a zoom-out effect as illustrated in Fig. 1(c), while taking $0 < \alpha < 1$ produces a zoom-in effect as illustrated in Fig. 1(d).

3.2. Spline Based Algorithm

In this section, we alternatively reconstruct the modified phase $\tilde{\varphi}_k$ by performing cubic spline-based integration independently over disjoint spatial windows that tile the FM function over the pixel lattice \mathcal{O} . This requires an auxiliary technique for estimating the phase initial conditions that are needed to start the integration along the boundaries of each window.

Since each AM-FM component contributes to the reconstructed image only where the AM function is nonzero, the integration windows used in the spline-based reconstruction algorithm are automatically defined for each component around areas where the value of the AM function is above an empirical threshold. Integration is then carried out independently within each window. Because the frequency modulation is generally smoother at points where the AM function exhibits greater values, the initial conditions needed for the integration of each window are estimated from points that coincide with local maxima of the AM function. Cubic spline integration is then performed starting at these points to estimate the modified phase function $\tilde{\varphi}_k(\mathbf{p})$.

We used this approach to design two new modulation domain FM filters. We refer to the first of these new filters as the *frequency flattening filter*. The effect of the signal processing is to produce horizontal ridges with frequencies that are representative of the frequency magnitudes that were present in the original image, but without the original orientation information. For all points in an FM function, this filter sets the horizontal components of the instantaneous frequency vectors to zero, while the vertical components of the instantaneous frequency vectors are set to the magnitude of the original frequency vectors. The action of this filter is illustrated in Fig. 2(b), where the filter was applied to the *chirp* image of Fig. 2(a), which is identical to the image in Fig. 1(a).

The second FM filter performs frequency rotation. This filter modifies points in the FM function by rotating the instantaneous frequency vectors by a specified angle. Results of applying this filter to the *chirp* image are given in Fig. 2(c). For this example, the instantaneous frequency vectors were rotated by $\pi/2$ rad at all points in the frequency modulations of all AM-FM components. The magnitudes of the rotated instantaneous frequency vectors were also modified to be inversely proportional to the magnitudes of the original frequencies.

Finally, in order to show the effects of the frequency rotation filter on a natural image, we also applied the filter to the *barbara* image, which is shown in Fig. 2(d). In this case, the frequency rotation filter was applied to achieve the perceptually motivated signal processing goal of changing the stripes on Barbara's pants from vertical to horizontal. Hence, the frequency rotation filter was applied only in a spatial region of interest (ROI) indicated by the black box overlaid on the image in Fig. 2(d). The filter was designed to rotate the instantaneous frequency vectors of all AM-FM components by $\pi/2$ rad, but only at pixels where the original instantaneous frequency vectors satisfied $-\pi/2 < \psi_k < 3\pi/16$ and $\pi/2 < r_k < 2\pi$, where $\psi_k = \arg \nabla \varphi_k$ is the orientation of the instantaneous frequency vector, $r_k = |\nabla \varphi_k|$ is the magnitude of the instantaneous frequency vector given in units of rad/pixel, $k \in [1, K]$, and K is the number of computed AM-FM components. The ROI from the FM filtered image is shown in Fig. 2(f), where it can be seen that the desired processing goal was achieved.

4. CONCLUSION

In this paper, we proposed two new algorithms to perform multi-component phase reconstruction subsequent to the application of modulation domain FM filters. For the first time, the new phase reconstruction algorithms were used to demonstrate modulation domain filters that directly modify the FM functions of an image to achieve predictable, perceptually motivated image processing goals that would be difficult or impossible to achieve with ordinary LTI spatial filters.

These two algorithms significantly broaden the scope of processing that can be achieved with modulation domain image filtering. Important problems that remain open include the development of a greater understanding of expected modulation domain filter behavior, as well as the development of improved reconstruction algorithms. In addition, a better understanding of the exact relationships between different filtered AM-FM components would be useful in defining improved high-fidelity image reconstruction techniques.

5. REFERENCES

- [1] M.S. Pattichis and A.C. Bovik, "Analyzing image structure by multidimensional frequency modulation," *IEEE Trans. Pattern Anal., Machine Intel.*, vol. 29, no. 5, pp. 753–766, May 2007.
- [2] P. Maragos and A. C. Bovik, "Image demodulation using multidimensional energy separation," *J. Opt. Soc. Amer. A*, vol. 12, no. 9, pp. 1867–1876, Sep. 1995.
- [3] J. P. Havlicek, D. S. Harding, and A. C. Bovik, "Multidimensional quasi-eigenfunction approximations and multicomponent AM-FM models," *IEEE Trans. Image Proc.*, vol. 9, no. 2, pp. 227–242, Feb. 2000.
- [4] A. C. Bovik, N. Gopal, T. Emmoth, and A. Restrepo, "Localized measurement of emergent image frequencies by Gabor wavelets," *IEEE Trans. Info. Theory*, vol. 38, no. 2, pp. 691–712, Mar. 1992.
- [5] J. P. Havlicek, P. C. Tay, and A. C. Bovik, "AM-FM image models: Fundamental techniques and emerging trends," in *Handbook of Image and Video Processing*, A. C. Bovik, Ed., pp. 377–395. Elsevier Academic Press, Burlington, MA, 2nd edition, 2005.
- [6] P. Maragos, J. F. Kaiser, and T. F. Quatieri, "Energy separation in signal modulations with applications to speech analysis," *IEEE Trans. Signal Proc.*, vol. 41, no. 10, pp. 3024–3051, Oct. 1993.
- [7] T. Ezzat, J. Bouvrie, and T. Poggio, "AM-FM demodulation of spectrograms using localized 2D max-Gabor analysis," in *Proc. IEEE Int'l. Conf. Acoust., Speech, Signal Process.*, Honolulu, HI, Apr. 15-20, 2007, vol. IV, pp. 1061–1064.
- [8] C. T. Nguyen, J. P. Havlicek, and M. Yeary, "Modulation domain template tracking," in *Proc. IEEE Int'l. Conf. Comput. Vision, Pattern Recog.*, Minneapolis, MN, Jun. 17-22, 2007, 8 pp.
- [9] R. S. Prakash and R. Aravind, "Modulation-domain particle filter for template tracking," in *Proc. Int'l. Conf. Pattern Recog.*, Tampa, FL, Dec. 8-11, 2008, pp. 1–4.
- [10] R. A. Sivley and J. P. Havlicek, "Perfect reconstruction AM-FM image models," in *Proc. IEEE Int'l. Conf. Image Proc.*, Atlanta, GA, Oct. 8-11, 2006, pp. 2125–2128.
- [11] C. T. Nguyen, R. A. Sivley, and J. P. Havlicek, "First results in perceptually-based AM-FM image filtering," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Santa Fe, NM, Mar. 24-26, 2008, pp. 77–80.
- [12] C. T. Nguyen and J. P. Havlicek, "AM-FM image filters," in *Proc. IEEE Int'l. Conf. Image Proc.*, San Diego, CA, Oct. 12-15, 2008, pp. 789–792.
- [13] R. A. Sivley and J. P. Havlicek, "Multidimensional phase unwrapping for consistent APF estimation," in *Proc. IEEE Int'l. Conf. Image Proc.*, Genoa, Italy, Sep. 11-14, 2005, vol. II, pp. 458–461.
- [14] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, "Shiftable multi-scale transform," *IEEE Trans. Info. Theory*, vol. 38, no. 2, pp. 587–607, March. 1992.
- [15] E. P. Simoncelli and W. T. Freeman, "The steerable pyramid: a flexible architecture for multi-scale derivative computation," in *Proc. IEEE Int'l. Conf. Image Proc.*, Washington, DC., Oct. 23-26, 1995, pp. 444–447.
- [16] R. A. Sivley and J. P. Havlicek, "A spline-based framework for perfect reconstruction AM-FM models," in *Proc. IEEE Southwest Symp. Image Anal., Interp.*, Denver, CO, Mar. 26-28 2006, pp. 198–202.
- [17] M. Unser, A. Aldroubi, and M. Eden, "B-spline signal processing: Part I-theory," *IEEE Trans. Signal Proc.*, vol. 41, no. 2, pp. 821–833, Feb. 1993.
- [18] M. Unser, A. Aldroubi, and M. Eden, "B-spline signal processing: Part II-efficient design and applications," *IEEE Trans. Signal Proc.*, vol. 41, no. 2, pp. 834–848, Feb. 1993.
- [19] D. C. Ghiglia and L. A. Romero, "Robust two-dimensional weighted and unweighted phase unwrapping that uses fast transforms and iterative methods," *J. Opt. Soc. Amer. A*, vol. 11, no. 1, pp. 107–117, Jan. 1994.




Fig. 1. FM filtering with least squares phase reconstruction. (a) Original chirp image. (b) Chirp image with rotated FM functions. (c) Isotropic zoom-out effect with $\alpha = 1.2$. (d) Anisotropic zooming: vertical zoom-in ($\alpha = 0.7$) combined with horizontal zoom-out ($\alpha = 1.2$).




Fig. 2. FM filtering with spline-based phase reconstruction. (a) Original radial chirp image. (b) Result of frequency flattening filter. (c) Result of frequency rotation and inversion filter. (d) Original 512x512 barbara image. The region of interest (ROI) is denoted by a black box. (e) Detail view of the ROI. (f) Selective frequency rotation within the ROI; the signal processing goal was to change the stripes on Barbara's pants from vertical to horizontal. Note that certain stripes on the shawl were also rotated since they fell within the ROI and matched the original orientation and granularity of the pants.