Stanford Exploration Project, Report 100, April 20, 1999, pages 211–223

Stanford Exploration Project, Report 100, April 20, 1999, pages 211–223

Texture synthesis and prediction error filtering

Morgan Brown[1]

keywords: *texture, PEF, simulation, helix, coherency, missing data*

ABSTRACT

The spectrum of a prediction-error filter (PEF) tends toward the inverse spectrum of the data from which it is estimated. I compute 2-D PEF's from known "training images" and use them to synthesize similar-looking textures from random numbers via helix deconvolution. Compared to a similar technique employing Fourier transforms, the PEF-based method is generally more flexible, due to its ability to handle missing data, a fact which I illustrate with an example. Applying PEF-based texture synthesis to a stacked 2-D seismic section, I note that the residual error in the PEF estimation forms the basis for "coherency" analysis by highlighting discontinuities in the data, and may also serve as a measure of the quality of a given migration velocity model. Last, I relate the notion of texture synthesis to missing data interpolation and show an example.

INTRODUCTION

In terms of digital images, the word *texture* might be defined as, "an attribute representing the spatial arrangement of gray levels of the pixels in a region," (IEEE, 1990). In the same context, I define *texture synthesis* as the process of first estimating the spatial statistical properties of a known image and then imparting these statistics onto a second (random) image. Figure 1 illustrates the general approach taken here: an uncorrelated image is transformed into one with the same statistical qualities as a known "training image" (TI), through an as-yet undefined filtering operation.

Texture synthesis is an active area of research in the computer graphics community, owing to the need for realistic, quickly generated surface textures (Simoncelli and Portilla, 1998; Heeger and Bergen, 1995; Brown and Mao, 1998), but the same notion of texture applies to the earth sciences as well. Physically measurable quantities, be they geology, gravity, or topography, behave in certain repeatable ways as a function of space, i.e., these quantities have a given texture. Inversion problems are often underdetermined, hampered by a lack of "hard" measurements, causing a nullspace of high dimension. A priori "soft" constraints on functional form of the unknown model help in suppressing the nullspace of modeling operators. These a priori constraints can be conceptualized as textures. For instance, in velocity analysis and tomography, the earth's velocity field is sometimes assumed to have a

[1]**email:** morgan@sep.stanford.edu

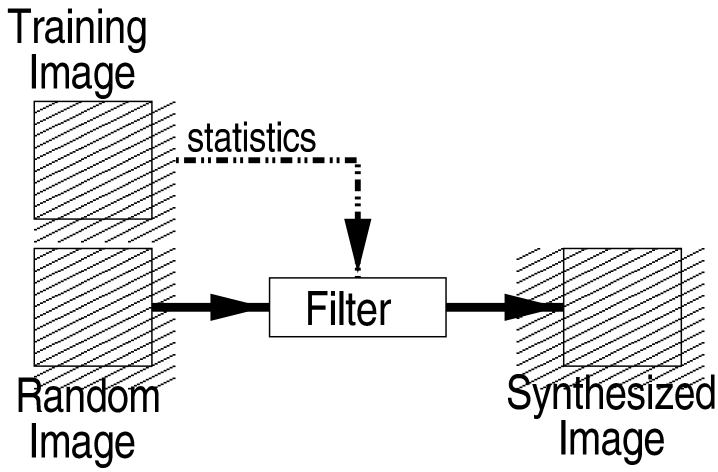

Figure 1: The generalized texture synthesis algorithm. From the training image (TI), statistics are extracted and encoded into a filtering operation, which forces an uncorrelated image to have the same statistical qualities, or *texture*, as the TI. morgan1-syn-templ [NR]

"blocky" texture (Clapp et al., 1998). Underdetermined inverse interpolation problems are often regularized by assuming "smooth" model texture (Claerbout, 1998).

The prediction-error filter (PEF) is an autoregressive filter which has the distinction of capturing the inverse spectrum of the data it is regressed upon. Because it captures this essential statistical property of the data, the PEF is a candidate for the generic "filter" operation shown in Figure 1.

This paper is intended as a follow-up to the earlier work by Claerbout and Brown (1999), which presented a texture synthesis technique utilizing 2-D PEF's and 2-D deconvolution via the helix transform (Claerbout, 1998b). First I motivate the texture synthesis problem by applying a Fourier transform-based technique to create synthetic textures of everyday objects, then introduce and apply a PEF-based technique to synthesize the same images. I compare the results of the two methods and conclude that the PEF-based method is the better choice because it more naturally handles missing data. Next I apply the PEF-based method to a 2-D stacked seismic section. The nature of the residual error in the PEF estimation of this example suggests application to seismic discontinuity detection and migration velocity analysis. Last, I solve a simple missing data problem to illustrate how regularization with a PEF imparts a reasonable "texture" onto the nullspace.

FOURIER TRANSFORM METHOD

The texture synthesis methodology of this paper really boils down to one of spectral estimation. An image's amplitude spectrum contains the relative weights between frequency components, while the phase spectrum localizes these frequency components in space (Castleman, 1996). Therefore, it stands to reason that the texture of stationary, loosely correlated images is adequately modeled using the amplitude spectrum alone. This idea is the basis for the Fourier Transform method of texture synthesis: all "realizations" of texture synthesis are forced to have the same amplitude spectrum, differing only in phase. The following is an outline of the method.

1. Given a training image, $t(x, y)$, compute its amplitude spectrum:

$$R(k_x, k_y) = T^*(k_x, k_y)T(k_x, k_y) \tag{1}$$

2. Create random phase function: $\phi_r(k_x, k_y) = random\ numbers.$

3. Reconstruct by substituting random phase:

$$t_{\text{recon}}(x, y) = \mathcal{F}^{-1}\left\{ \sqrt{\left|R(k_x, k_y)\right|}\ e^{i\phi_r(k_x, k_y)} \right\} \tag{2}$$

Figures 2 and 3 illustrate the Fourier transform method of texture synthesis. Clockwise from top-left: the training image, the synthesized image, the TI's amplitude spectrum, and the TI's phase spectrum.

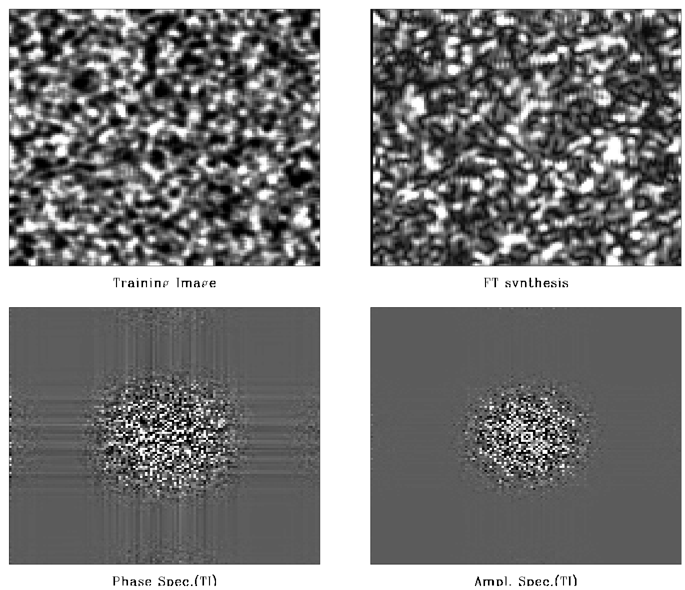

Figure 2: Smoothed random image and Fourier transform synthesis. The TI is stationary, so the synthesis result is convincing. Notice that the true phase, in the regions where the modulating amplitude spectrum is nonzero, is quite random in appearence. morgan1-rand2d-ftsyn [ER]

PEF-BASED METHOD

Theoretically, the convolution of data (N_d points) and a PEF (N_a coefficients) estimated from the data is approximately uncorrelated in the limit $N_a \to N_d \to \infty$: a spike at zero lag plus Gaussian, independent identically distributed (iid) noise elsewhere. Thus the spectrum of this residual error is approximately white. The frequency response of the "inverse PEF", as computed by deconvolution, is an N_a-point parameterization of the N_d-point inverse amplitude spectrum, as illustrated in Figure 4. As the size of the filter increases, the parameterization becomes more accurate, as expected from theory (Claerbout, 1976). The notion of PEF as "decorrelator" is quite akin to decomposition by principal components (Castleman, 1996), where the number of principal components used in computation determines the degree of decorrelation.

The following is an outline of the PEF-based texture synthesis method.

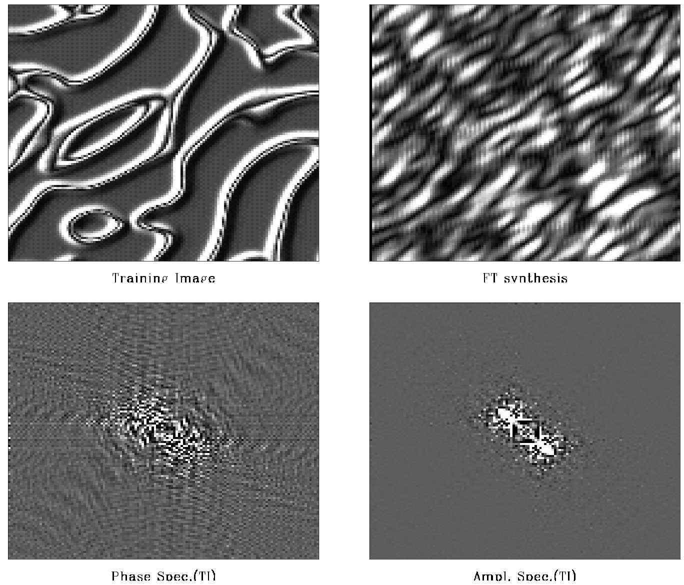

Figure 3: "Ridges" image and Fourier transform synthesis. The correlation is both long-range and extremely complicated - quite like a meandering network of fluvial channels. Though the synthesized image has the same *general* character as the TI, not all of the structures are modeled, proving the inadequacy of the amplitude spectrum for modeling nonstationary, highly correlated images. The TI phase spectrum shows some ordering, so the random phase substitution was ill-advised. morgan1-ridges-ftsyn [ER]

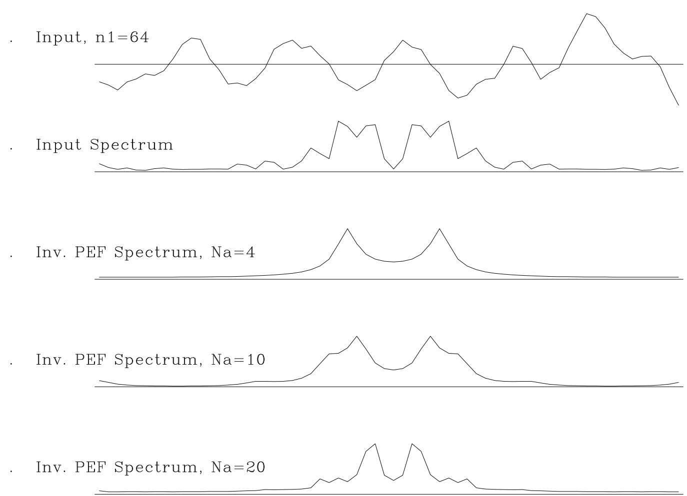

Figure 4: Frequency response of “inverse PEF” (deconvolution) as a function of filter size. As expected, as the filter length increases, the approximation improves. morgan1-rand1d-spec [ER]

1. Given training image $t(x,y)$, estimate unknown PEF $a(x,y)$ via least squares minimization:

$$\min \| t * a \|^2 \tag{3}$$

2. The residual $r = t * a$ is approximately uncorrelated, with the same dimension as the TI, since we use an "internal" convolution algorithm (Claerbout, 1998). It can be proved that a is a minimum phase filter, (Claerbout, 1976) so deconvolution (polynomial division) robustly and stably reconstructs t given r. Generate a random residual r' with the same dimension as r. To create the synthetic texture, simply deconvolve r' by a:

$$t_{\text{syn}} = r'/a \tag{4}$$

   where the “ / ” refers to polynomial division, our preferred method of deconvolution.

Though the residual is uncorrelated, it does contain “phase” information. Deconvolution of a random image blindly spreads scaled copies of the impulse response of the inverse PEF across the output space. If the residual r is not sufficiently whitened, then the replacement of r with r' will lead to an ineffective representation of t by t_{syn}.

Figures 5 through 7 illustrate the PEF-based texture synthesis process. The left-hand panel shows the training image, the center panel shows the residual $r = t * a$, and the right-hand panel shows the synthesized image, $t_{\text{syn}} = r'/a$. A 10x10 PEF is used in each case. The blank areas in the residual panel correspond to regions where the PEF falls outside the bounds of the known data.

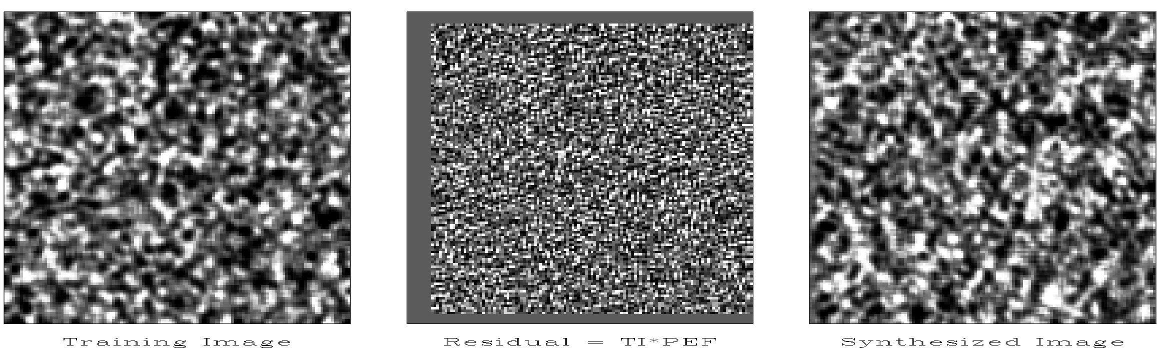

Figure 5: Smoothed random 2-D image and PEF-based texture synthesis result. The TI is quite simple (stationary, low correlation), so as expected, the synthesized image and the TI are almost indistinguishable. To the naked eye, the residual appears effectively white. morgan1-rand2d-pefsyn [ER]

WHY USE THE PEF?

PEF-based texture synthesis can only achieve the results of the Fourier transform method (Figures 2 and 3) in the limit $N_a \rightarrow N_d$, which is unrealistic in practical situations, where N_d is very large. Least squares estimation of the filter in this case is certainly costlier than

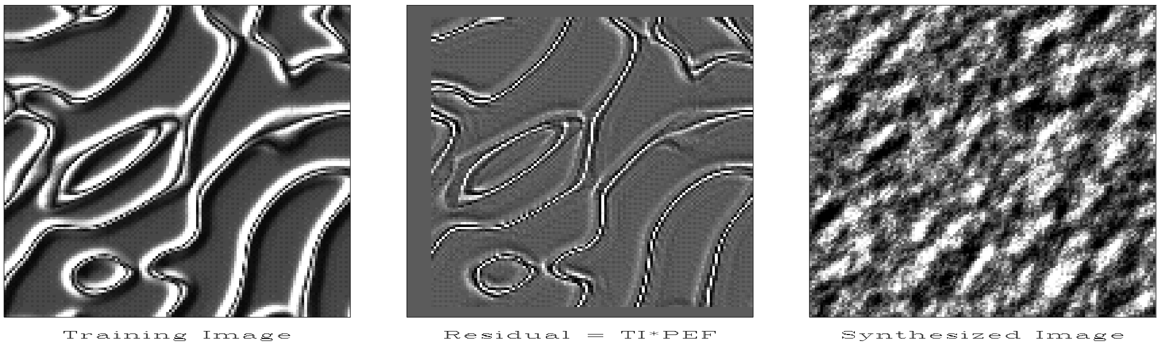

Figure 6: "Ridges" image and PEF-based texture synthesis result. Recall that the complicated connected features of this image were not completely synthesized by the Fourier transform method (Figure 3), of which the PEF method is an approximation. This synthesized image bears even less resemblance to the TI, exhibiting only a general southwest-to-northeast trend. The wavy, ridge-like features have many different dips, making them difficult to predict with a PEF, and with two point statistics in general. The same can be said for the ubiquitous hyperbolic features of reflection seismology. morgan1-ridges-pefsyn [ER]

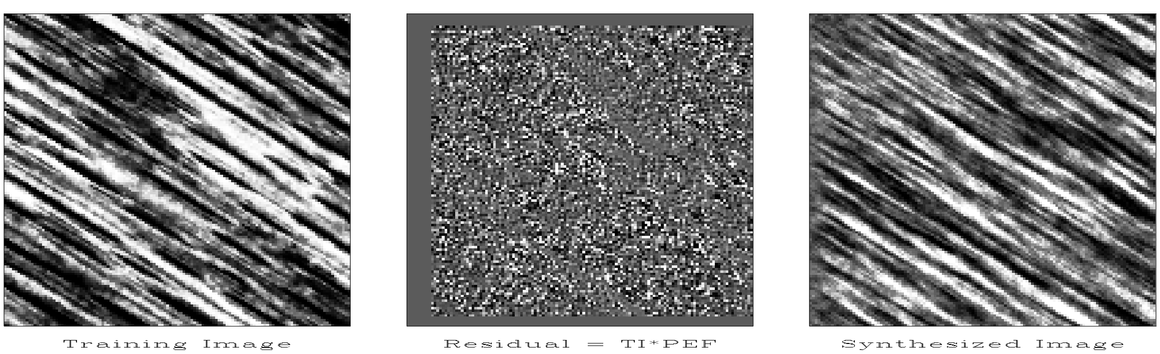

Figure 7: "Wood" image and PEF-based texture synthesis result. The synthesis result is pleasing. The PEF-based method preserves the general trend and relative scale length of the lineations in the TI. The correlation of the TI is relatively long-range, in that the lineations cross a large portion of the image, but the features are merely straight lines at one dip. morgan1-wood-pefsyn [ER]

three Fast Fourier transforms. On the other hand, if the filter size can be limited without compromising quality, which is the case for stationary, simply correlated images, then the PEF-based method is more flexible. Unlike the Fourier transform a PEF can be estimated easily when data are missing. Figure 8 shows that the PEF estimated from the incomplete data captures enough features of the data's spectrum to make a fairly convincing texture synthesis result. The output of the PEF-based method can be of any size, while the output of a Fourier transform is generally constrained to be the same size as the input.

APPLICATIONS

PEF Estimation with incomplete data

Modern reservoir characterization efforts take a pragmatic view of collected data. Rather than wait for collection of the elusive "perfect" dataset, the desire is to incorporate a wide variety of possibly incomplete data types into a single inversion scheme (Caers and Journel, 1998). Often the only data available is spatially incomplete. Figure 8 shows the result of texture synthesis on training images with large void regions. As noted earlier, the blank areas in the center panels of the figure correspond to regions where the filter can't fit without falling on one or more missing points. Each of the "in-bounds" data points contributes one equation to the LS estimation of the 100 or so filter coefficients. Even when well over half of the data points are removed from the training image this result shows that we can still safely estimate a filter and synthesis a believable texture.

2-D Stacked Seismic Section

Figure 9 shows the result of applying PEF texture synthesis to a 2-D stacked seismic section. The residual panel is interesting; notice uncollapsed diffraction hyperbolas, two highlighted fault planes, and also statics-like artifacts in the earlier times. PEF's easily predict straight lines (plane waves) and sinusoids, but hyperbolas and discontinuities are quite another matter.

Matthias Schwab used the "plane wave prediction" property of the PEF in his Ph.D. thesis (Schwab, 1998) to create so-called "coherency cubes" from 3-D seismic data by nonstationary convolution with small PEF's. Development of viable seismic coherency attributes merits considerable industrial interest, as evidenced by the concentration of related articles in the March, 1999 edition of *The Leading Edge*.

If a good velocity model is used, poststack migration should collapse these hyperbolas, so one measure of the fitness of a given velocity model could be the relative amount of residual energy in the data*PEF panel. Additionally, to the same end, this technique could be used to measure the relative amount of residual curvature in common reflection point (CRP) gathers, which are flattened when the correct migration velocity is used (Biondi, 1997). This preprocessing could be done quickly, for the necessary PEF's are small.

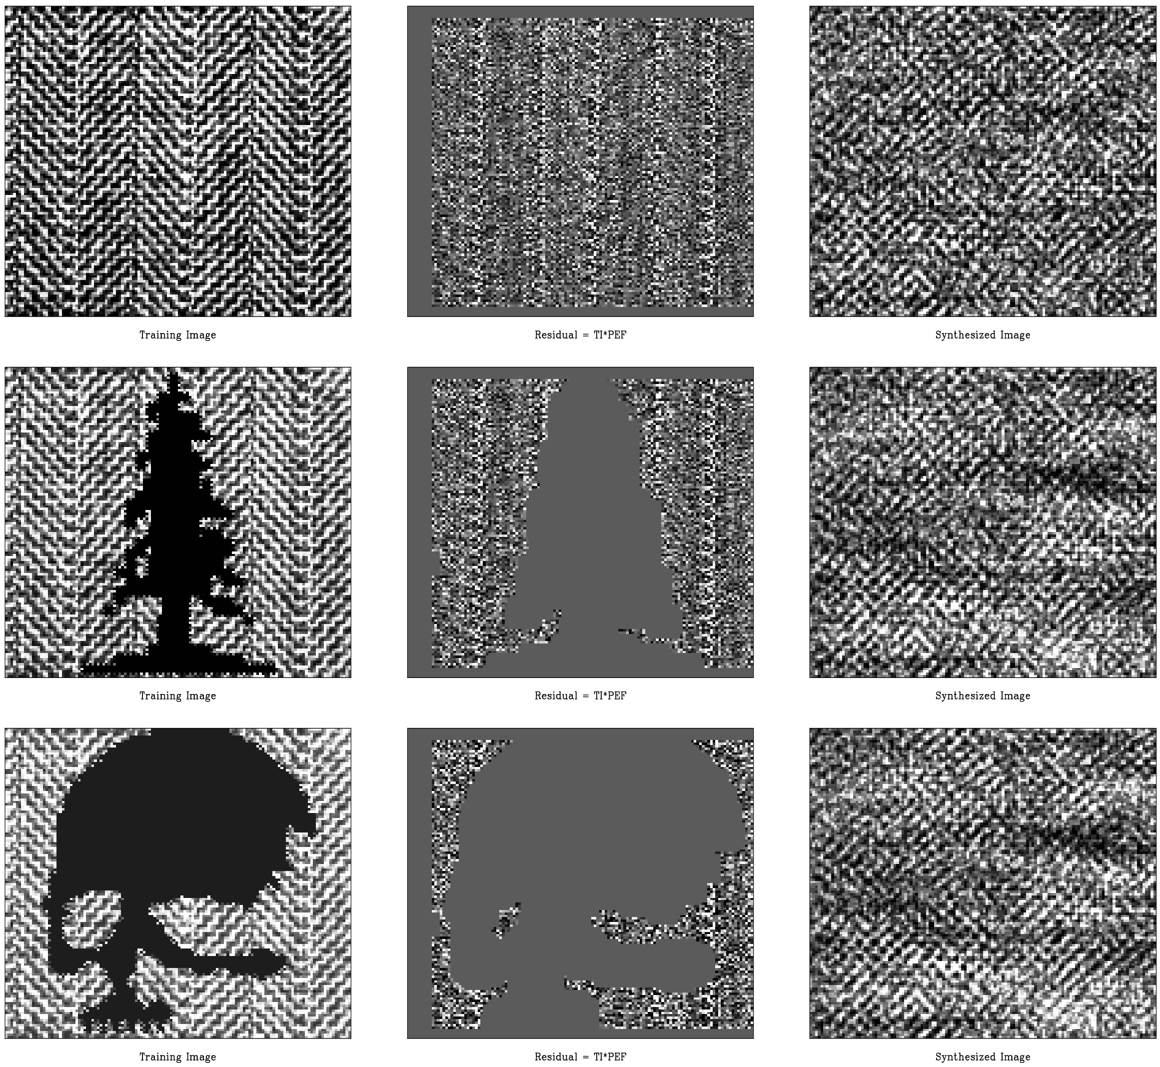

Figure 8: Comparison showing the effects of missing data on the PEF texture synthesis result, for two different "holes". Although half or more of the equations are removed from the PEF estimation problem, the synthesized textures still capture the character of the training image. Fourier transforms are ill-defined on irregular coordinate systems, but the PEF makes an estimate of the known data's spectrum regardless. morgan1-holes [ER]

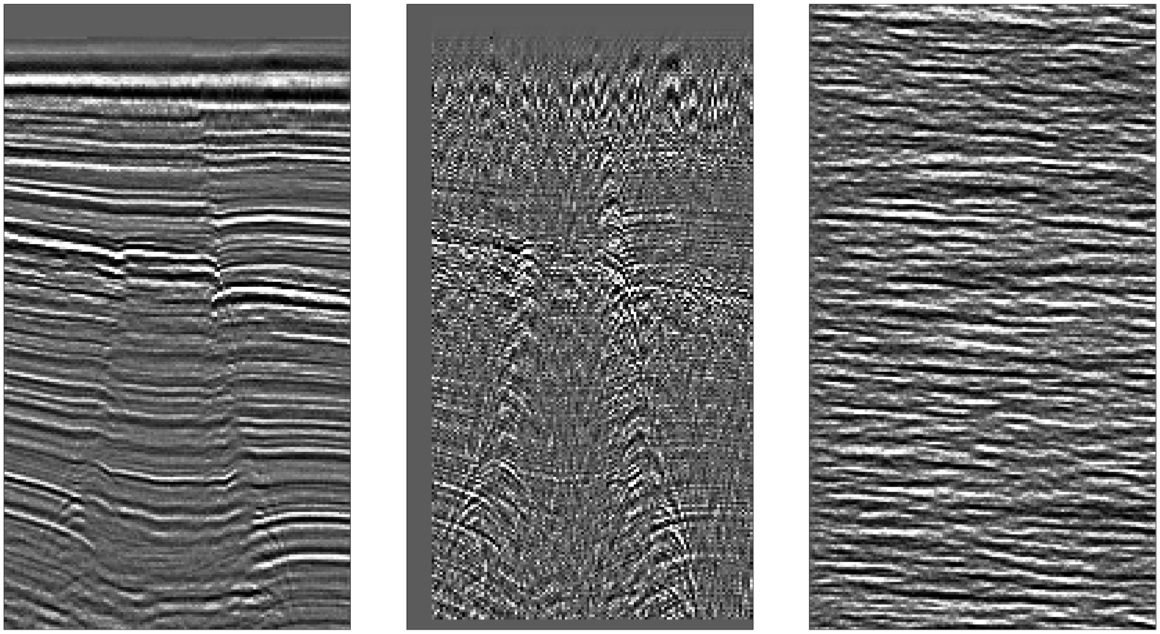

Figure 9: Stacked 2-D seismic section. morgan1-WGstack-pefsyn [ER]

Preconditioned Missing Data Infill

To fill "holes" in collected data, we have the familiar SEP formulation (Claerbout, 1998):

$$\mathbf{Km} - \mathbf{d} \approx 0 \tag{5}$$

$$\epsilon \mathbf{Am} \approx 0 \tag{6}$$

[5] is the "data matching" goal, which states that the model $\mathbf{m}$ must match the known data $\mathbf{d}$, while [6] is the "model smoothness" goal, where $\mathbf{A}$ is an arbitrary roughening operator. To combat slow convergence, Claerbout (1998) preconditions with the inverse of the convolutional operator $\mathbf{A}$ (multidimensional *de*convolution). Provided that $\mathbf{A}$ is minimum phase or factorizable into the product of minimum phase filters (Sava et al., 1998), the helix transform now permits stable multidimensional deconvolution. Making the change of variables $\mathbf{m} = \mathbf{A}^{-1}\mathbf{x}$, we have the equivalent preconditioned problem:

$$\mathbf{KA}^{-1}\mathbf{x} - \mathbf{d} \approx 0 \tag{7}$$

$$\epsilon \mathbf{x} \approx 0 \tag{8}$$

The operator $\mathbf{K}$ effectively maps vectors in model space into a smaller-dimension "known data space", so it has a nonempty nullspace. Missing points in model space are completely unconstrained by $\mathbf{K}$, so our choice of $\mathbf{A}$ wholly determines the behavior of the missing model points, i.e., their *texture* (Fomel et al., 1997). The PEF is a perfect choice for $\mathbf{A}$, as shown in Figure 10. The preconditioned, PEF-regularized result fills the hole quite believably after only 20 iterations, as opposed to the case where $\mathbf{A} = \nabla^2$, which imposes an unrealistically smooth texture on the missing model points.

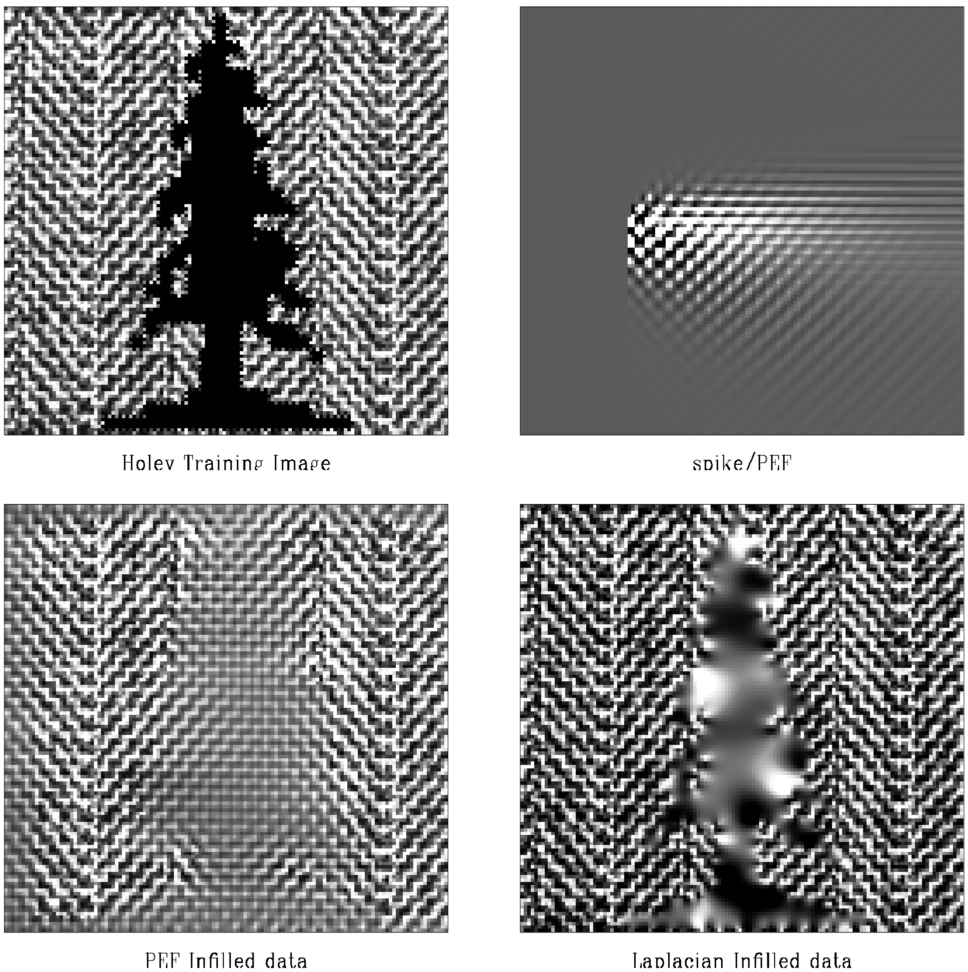

Figure 10: Clockwise from top left: Data with hole, impulse response of "inverse PEF" (deconvolution of the PEF estimated from the data and a spike), data in-filled using ∇^2 regularization, data in-filled using preconditioned PEF regularization. morgan1-tree-hole-filled [ER]

DISCUSSION

The goal of this paper is not to make slick surface textures for computer games. Nontheless, as a tutorial device, texture synthesis using the PEF is valuble, since it concretely and intuitively illustrates in two dimensions some of the fundamental concepts of autoregression which are proved only in the one dimensional case (Claerbout, 1976). In fact, some of the results shown here and in Claerbout and Brown (1999) have recently been incorporated into Jon Claerbout's textbook, *Geophysical Estimation by Example* (1998).

Both the Fourier transform and PEF-based texture synthesis operate under the assumption that the training image is sufficiently well characterized by amplitude spectrum alone. For some images (Figures 2, 5, and 7) the assumption holds, but for others (Figures 3, 6) it is obviously violated. Real digital images and earth phenomena alike often exhibit complex spatial correlation which are modelable only with multiple point templates (Caers and Journel, 1998; Malzbender and Spach, 1993). Additionally, I have ignored the interesting subjects of nonstationarity and spatial scale variance. By scale-variant, I mean that the characteristic scale of an image's features is not constant with respect to spatial frequency. Many methods for characterizing scale-variant images appeal to the world of wavelets for a methodology known as *multiresolution analysis* (Simoncelli and Portilla, 1998; Heeger and Bergen, 1995; Strang and Nguyen, 1997). The notion of texture synthesis for nonstationary images is ill-defined, since it amounts to a random reordering of filters estimated on locally-stationary patches, followed by deconvolution on the correspondng patches.

When the training image has missing values, as in Figure 8, the PEF-based texture synthesis method performs favorably. As shown in the missing data interpolation example (Figure 10), the ability of the PEF to reliably estimate the data spectrum, even with missing data, makes it an ideal regularization operator. Figure 9 illustrates the fact that the PEF primarily predicts plane waves. I proposed using a PEF residual measure to determine the viability of a given migration velocity. In general, PEF estimation/convolution might have value as a preprocessing step for a variety of applications. For instance, a very small PEF (2 columns) has a relatively large residual in the presence of conflicting dips, and thus may help in determining local filter size or patch size.

ACKNOWLEDGEMENTS

All the results in this paper were generated quite easily using Sergey Fomel's helix inversion library. All important programs have been included in the most recent release of SEPLib. The programs are conducive to curious experimentation, so I encourage the reader to use the makefile which accompanies the source code for this paper as a template.

REFERENCES

Biondi, B. 3-d seismic imaging:. http://sepwww.stanford.edu/sep/biondo/Lectures, 1997.

Brown, M., and Mao, S. Psych 267 final project - texture synthesis:. http://sepwww.stanford.edu/sep/morgan/texturematch/, 1998.

Caers, J., and Journel, A. G., 1998, Stochastic reservoir modeling using neural networks trained on outcrop data.: SPE Annual Technical Conference and Exhibition, 49026.

Castleman, K. R., 1996, Digital image processing: Prentice-Hall, Inc., Upper Saddle River, NJ.

Claerbout, J., and Brown, M., 1999, Two-dimensional textures and prediction-error filters: 61st Ann. Mtg., Eur. Assoc. Expl. Geophys.

Claerbout, J. F., 1976, Fundamentals of Geophysical Data Processing: Blackwell.

Claerbout, J. Geophysical Estimation by Example: Environmental soundings image enhancement:. http://sepwww.stanford.edu/sep/prof/, 1998.

Claerbout, J. F., 1998b, Multi-dimensional recursive filtering via the helix: Geophysics, **63**, no. 5, 1532–1541.

Clapp, R. G., Sava, P., and Claerbout, J. F., 1998, Interval velocity estimation with a null-space: SEP–**97**, 147–156.

Fomel, S., Clapp, R., and Claerbout, J., 1997, Missing data interpolation by recursive filter preconditioning: SEP–**95**, 15–25.

Heeger, D. J., and Bergen, J. R., 1995, Pyramid-based texture analysis/synthesis: Computer Graphics Proceedings, 229–238.

IEEE, 1990, IEEE standard glossary of image processing and pattern recognition terminology, standard 610.4-1990: IEEE Press, New York.

Malzbender, T., and Spach, S., 1993, A context sensitive texture nib: Proceedings of Computer Graphics International, 151–163.

Sava, P., Rickett, J., Fomel, S., and Claerbout, J., 1998, Wilson-Burg spectral factorization with application to helix filtering: SEP–**97**, 343–352.

Schwab, M., 1998, Enhancement of discontinuities in seismic 3-D images using a Java estimation library: Ph.D. thesis, Stanford University.

Simoncelli, E., and Portilla, J., 1998, Texture characterization via joint statistics of wavelet coefficient magnitudes: 5th IEEE Int'l Conf on Image Processing.

Strang, G., and Nguyen, T., 1997, Wavelets and Filter Banks: Wellesley-Cambridge Press.