THE DISCRETE ORTHONORMAL STOCKWELL TRANSFORM FOR IMAGE RESTORATION

Yanwei Wang

Department of Applied Mathematics
University of Waterloo,
200 University Avenue West,
Waterloo, ON, Canada, N2L 3G1

Jeff Orchard

David R. Cheriton School of Computer Science
200 University Avenue West,
University of Waterloo,
Waterloo, ON, Canada, N2L 3G1

ABSTRACT

This paper describes an automated image restoration algorithm. The technique is based on the Stockwell transform (ST) and its discrete version, the discrete orthonormal Stockwell transform (DOST). These mathematical transforms provide a multiresolution spatial-frequency representation of a signal or image.

First, we give a brief introduction to the Stockwell transform, and the DOST. Then we describe a restoration method using the DOST based on the total variation (TV) minimization model. The results show that the DOST restoration outperforms the wavelet restoration by giving a higher Peak Signal to Noise Ratio (PSNR).

Index Terms— DOST, Stockwell transform, image restoration, inpainting.

1. INTRODUCTION

Images have been compressed using the wavelets since the 1970s. When the compressed images are transferred over the wired or wireless medias, the wavelet coefficients are transferred. Due to the transportation, the routers and many other reasons, it is possible to get partially lost on the the coefficients. To restore the partially lost information we need to use restoration algorithms to recover the images to the level acceptable to human eyes.

Even though the wavelets have dominated image compression for years, other competitive methods for compression are being worked on. In Ref. [1] , an compression algorithm was described using the newly invented discrete orthonormal Stockwell transform (DOST) [2]. The DOST compression outperformed the wavelet compression in the base level and using the DOST, new technologies are expected to be comparable to a higher level of wavelets compression. The possibility to restore a DOST-compressed image with partially lost information (coefficients) will make the reality of DOST compression more convincing.

2. REVIEW OF THE ST AND THE DOST

The Stockwell transform (ST), proposed in 1996 [3, 4, 5, 6], gives a full time-frequency decomposition of a signal. Consider a one-dimensional function $h(t)$. The Stockwell transform of $h(t)$ is defined as the Fourier transform (FT) of the product between $h(t)$ and a Gaussian window function,

$$S(\tau, f) = \int_{-\infty}^{\infty} h(t) \frac{|f|}{\sqrt{2\pi}} e^{-\frac{(\tau-t)^2 f^2}{2}} e^{-i2\pi f t} dt, \quad (1)$$

where f is the frequency, and t and τ are time variables. In this way, the ST decomposes a signal into temporal (τ) and frequency (f) components.

By the integral properties of the Gaussian function, the relation between $S(\tau, f)$ and $H(f)$ (the Fourier transform of $h(t)$) is

$$\int_{-\infty}^{\infty} S(\tau, f) d\tau = H(f). \quad (2)$$

That is, the accumulation of the Stockwell coefficients over the time axis yields the FT of the signal, highlighting a special feature of the ST. Hence, the original function $h(t)$ can be recovered by calculating the inverse Fourier transform of $H(f)$,

$$h(t) = \int_{-\infty}^{\infty} \left\{ \int_{-\infty}^{\infty} S(\tau, f) d\tau \right\} e^{i2\pi f t} df. \quad (3)$$

Using the equivalent frequency-domain definition of the Stockwell transform, the Discrete Stockwell transform (DST) [3] can be written

$$S[j, n] = \sum_{m=0}^{N-1} H[m+n] e^{-2\pi^2 m^2/n^2} e^{i2\pi m j/N}, \quad (4)$$

for $n \neq 0$, where $H[\cdot]$ is the DFT of $h[\cdot]$. For the $n = 0$ voice, define

$$S[j, 0] = \frac{1}{N} \sum_{m=0}^{N-1} h[m], \quad (5)$$

analogous to the DC value of the FT.

The DST has been used in various fields [2], but from (4), it is obvious that the ST is an overcomplete representation. For a signal of length N, there are N^2 Stockwell coefficients, and each one takes $\mathcal{O}(N)$ to compute. Hence, computing all N^2 coefficients of the ST has computational complexity $\mathcal{O}(N^3)$. The ST gets exponentially more expensive for higher-dimensional. A more efficient mathematical and computation framework was needed to pursue this time-frequency decomposition.

The DOST is a pared-down version of the fully redundant ST [2]. Since lower frequencies have longer periods, it stands to reason that lower frequencies can cope with lower sampling rates. Hence, the DOST subsamples the low frequencies. Similarly, high frequencies have higher sampling rates. The DOST takes advantage of this sample spacing paradigm, and distributes its coefficients accordingly. It does so by constructing a set of N orthogonal unit-length basis vectors, each of which targets a particular region in the time-frequency domain. Which region is dictated by a set of parameters: ν specifies the center of each frequency band (voice), β is the width of that band, and τ specifies the location in time. Using these parameters, the kth basis vector is defined as

$$D[k]_{[\nu,\beta,\tau]} = \frac{1}{\sqrt{\beta}} \sum_{f=\nu-\beta/2}^{\nu+\beta/2-1} \exp\left(-i2\pi\frac{k}{N}f\right) \times \exp\left(i2\pi\frac{\tau}{\beta}f\right) \exp\left(-i\pi\tau\right), \quad (6)$$

for $k = 0, \cdots, N-1$, which can be summed analytically to

$$D[k]_{[\nu,\beta,\tau]} = ie^{-i\pi\tau}\frac{e^{-i2\alpha(\nu-\beta/2-1/2)} - e^{-i2\alpha(\nu+\beta/2-1/2)}}{2\sqrt{\beta}\sin\alpha}, \quad (7)$$

where $\alpha = \pi(k/N - \tau/\beta)$ is the center of the temporal window.

To make the family of basis vectors in (7) orthogonal, the parameters ν, β and τ have to be chosen suitably [1, 2]. For real applications, it is helpful to order these N coefficients into a 1-D vector. The ordering we use is shown in Fig. 1 for a signal of length 16. By convention, our time index (τ) traverses the time axis in the negative direction for negative frequencies. Doing so creates a symmetric correspondence between the positive- and negative-frequency coefficients in the 1-D representation. That is, for a given coefficient with index i in the 1-D DOST vector, its negative-frequency analog is at index $N - i$.

Based on (6), the DOST decomposition and reconstruction can be computed faster than the direct matrix multiplication. Though we do not report on it here, we have developed a way to speed up the DOST calculation to achieve a complexity of $\mathcal{O}(N \log N)$ for a signal of length N, and $\mathcal{O}(N^2 \log N)$ for an $N \times N$ image.

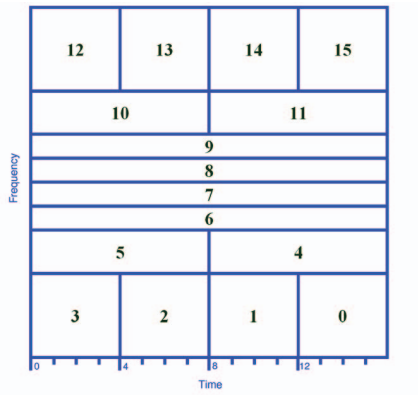

Fig. 1. The order of the 2-D DOST coefficients into an 1-D N-vector.

Fig. 2. Lena and the logarithm of its DOST coefficients.

Figure 2 shows the logarithm of the magnitude of the 2-D DOST coefficients for one of the most popular example images, Lena. The ST and the DOST are both separable transforms, so can be applied in higher dimensions trivially. As we can see, the coefficients decay very quickly, which makes the DOST a powerful tool for image compression. Moreover, the self-similarity over the decomposed spatial-frequency domain makes the DOST excellent for image inpainting and restoration.

3. METHOD AND ALGORITHM

The total variation (TV) minimization model has been used in denoising to get promising quality [7] and has started to benefit the image restoration. In Ref. [8], Chan et al. presented a method to restore a wavelet-compressed image by minimizing TV. Here, we follow their strategies and replace the wavelet coefficients with DOST coefficients.

Assume the original image is $u_0(x)$ and has the following DOST decomposition,

$$u_0(d, x) = \sum_{i,j} f_{i,j} D[x_1]_i D[x_2]_j, \quad (8)$$

where $f_{i,j}$ is the $(i,j)_{th}$ DOST coefficient (for simplicity of notation, the DOST coefficients for an image have

two indices, hence form a 2-D matrix as a whole) and $x = (x_1, x_2) \in \mathbf{R}^2$ is the 2-D spatial coordinate.

Model:

Define the total variation of the image as [8] [9]

$$F(u;d) = \int_{\mathbf{R}^2} |\nabla_x u(d,x)| dx = \text{TV}(u(d,x)), \qquad (9)$$

where $u(d,x)$ is the damaged image and has the DOST decomposition:

$$u(d,x) = \sum_{i,j} d_{i,j} D[x_1]_i D[x_2]_j, \qquad i,j \in \mathbf{Z}, \qquad (10)$$

with the constraint $d_{i,j} = f_{i,j}$, $(i,j) \notin \mathbf{I}$, where $\mathbf{I}$ is the inpainting index region.

To minimize the cost function $F(u;d)$, we let the first derivative of $F(u;d)$ – with respect to the coefficients $d_{i,j}$ – equal zero.

Algorithm:

1. Start with $n = 0$ and initial guess $d_{i,j}^{new} = d_{i,j}$. Set $d_{i,j}^{old} = 0$, and the initial error $E = ||d^{new} - d^{old}||_2$.
2. while $n < N$ or $E > \delta$, do
   - Set $d^{old} = d^{new}$,
   - Calculate d^{TV} as described after the main pseudo-code here.
   - For all (i,j), update

$$d_{i,j}^{new} = d_{i,j}^{old} + \frac{\Delta t}{\Delta x} \lambda_{i,j}, \qquad (11)$$

   where $\lambda_{i,j} = d_{i,j}^{TV} \chi_{i,j}$, and

$$\chi_{i,j} = \begin{cases} 1 & \text{if } (i,j) \in \mathbf{I} \\ 0 & \text{if } (i,j) \notin \mathbf{I} \end{cases}$$

   - Compute error $E = ||d^{new} - d^{old}||_2$, and set $n = n + 1$.

Inside the algorithm above the d^{TV}, which corresponds to the non-linear integral inside the minimization process can be calculated using $u = \text{IDOST}(d)$, where IDOST is the inverse DOST transform. For all (i,j), compute

$$\begin{aligned} \text{curv}_{i,j} &= D_1^- \left(\frac{D_1^+ u_{i,j}}{\sqrt{|D_1^+ u_{i,j}|^2 + |D_2^+ u_{i,j}|^2 + \epsilon}} \right) \\ &\quad + D_2^- \left(\frac{D_2^+ u_{i,j}}{\sqrt{|D_1^+ u_{i,j}|^2 + |D_2^+ u_{i,j}|^2 + \epsilon}} \right) \end{aligned} \qquad (12)$$

where $D_1^+ u_{i,j} = u_{i+1,j} - u_{i,j}$, $D_2^+ u_{i,j} = u_{i,j+1} - u_{i,j}$ are the forward differences, and $D_1^- u_{i,j} = u_{i,j} - u_{i-1,j}$,$D_2^- u_{i,j} = u_{i,j} - u_{i,j-1}$ are the backward differences, and ϵ is a small positive number that is used to prevent the numerical blow-up.

Then the projection of the curvature on the DOST basis can be calculated by $d^{TV} = \text{DOST(curv)}$.

4. RESULTS

Fig. 3. The original synthetic image for restoration test.

For our tests, we used the synthetic image shown in Fig. 3. To gain a valuable comparison, we also implemented the wavelet restoration method for wavelet-compressed images. The results from the numerical calculation are shown in Fig. 4 and Fig. 5.

As we can observe from Fig. 4 and Fig. 5, the DOST restoration has successfully restored more information than the wavelet restoration by offering a higher PSNR and more visible features and edges of the image. In the extreme case of randomly losing 90% of the coefficients, even though the original damaged image can hardly be recognized, the DOST restoration method still recovers some of the features and edges.

Based on the above experiment, another fact can also get confirmed. When the two methods are faced with equivalently-degraded images (having lost the same coefficients), the DOST-compressed image recovers more information than the wavelet-compressed image, suggesting that DOST compression is more resilient to packet loss.

5. CONCLUSION AND FUTURE WORK

The DOST restoration algorithm has outperformed the wavelet restoration over the corresponding compression techniques. We conjecture that even better performance can be realized using a more sophisticated restoration method that utilizes the self-similarity in the DOST coefficients.

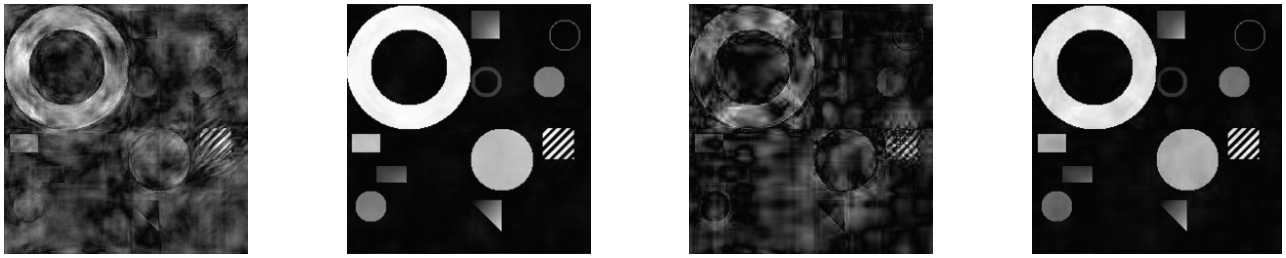

(a) DOST (b) DOST restored (c) Wavelets (d) Wavelets restored

Fig. 4. Image restoration test for randomly losing 50% of the DOST and wavelet coefficients. (a) is the damaged DOST compressed image with PSNR=11.17 and (c) is the damaged wavelet compressed image with PSNR=10.15. (b) (PSNR=28.94) is restored using the DOST. (d) (PSNR=26.28) is restored using the wavelets. As we can see, the DOST-restored image is also sharper and clearer than the wavelet restored image.

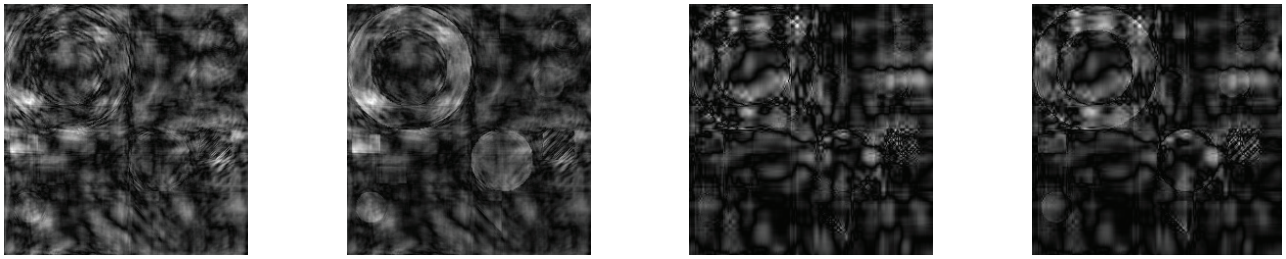

(a) DOST (b) DOST restored (c) Wavelets (d)Wavelets restored

Fig. 5. Image restoration test for randomly losing 90% of the DOST and wavelet coefficients. (a) is the damaged DOST compressed image with PSNR=8.83 and (c) is the damaged wavelet compressed image with PSNR=8.75. (b) (PSNR=9.06) is restored using the DOST. (d) (PSNR=8.93) is restored using the wavelets. This is an extreme test of the restoration, with heavy loss of information in both images. Even though neither method can restore all the features and edges, the DOST restoration method restores more visible image characteristics than the wavelet restoration.

6. REFERENCES

[1] Yanwei Wang and Jeff Orchard, "On the use of the Stockwell transform for image compression," *IS&T/SPIE Electronic Imaging*, January 2009, 7 pages, accepted on September 19, 2008, to appear recently.

[2] R. G. Stockwell, "A basis for efficient representation of the S-transform," *Digital Signal Processing*, vol. 17, no. 1, pp. 371–393, January 2007.

[3] R. G. Stockwell, L. Mansinha, and R. P. Lowe, "Localization of complex spectrum: the S transform," *IEEE Transactions on Signal Processing*, vol. 144, no. 4, pp. 998–1001, April 1996.

[4] L. Mansinha, R. G. Stockwell, and R. P. Lowe, "Pattern analysis with two-dimensional spectral localisation: Application of two-dimentional S transforms," *Physica A*, vol. 239, pp. 286–295, 1997.

[5] L. Mansinha, R. G. Stockwell, R.P. Lowe, M. Eramian, and R. A. Schincariol, "Local S-spectrum analysis of 1-D and 2-D data," *Phys. Earth Planet. Inter.*, vol. 103, pp. 329C336, November 1997.

[6] M. Eramian, R. A. Schincariol, Robert G. Stockwell, Robert P. Lowe, and L. Mansinha, "Review of applications of 1D and 2D S-transforms," *Proc. SPIE*, vol. 3078, pp. 558–568, April 1997.

[7] C. R. Vogel and M. E. Oman, "Iterative methods for total variation denoising," *SIAM J. Sci. Comput*, vol. 17, pp. 227–238, 1996.

[8] Tony F. Chan, Jianhong Shen, and Hao min Zhou, "Total variation wavelet inpainting," *J. Math. Imaging Vision*, vol. 25, pp. 107–125, 2006.

[9] Tony F. Chan, Jianhong Shen, and Hao min Zhou, "A total variation wavelet inpainting model with multilevel fitting parameters," *Proc. SPIE*, vol. 6313, 63130C, 2006.