

On the Use of the Stockwell Transform for Image Compression

Yanwei Wang^a and Jeff Orchard^b

^aDepartment of Applied Mathematics;

^bDavid R. Cheriton School of Computer Science,
200 University Avenue West,

University of Waterloo,
Waterloo, Ontario, Canada N2L 3G1

ABSTRACT

In this paper, we investigate the use of the Stockwell Transform for image compression. The proposed technique uses the Discrete Orthogonal Stockwell Transform (DOST), an orthogonal version of the Discrete Stockwell Transform (DST). These mathematical transforms provide a multiresolution spatial-frequency representation of a signal or image.

First, we give a brief introduction for the Stockwell transform and the DOST. Then we outline a simplistic compression method based on setting the smallest coefficients to zero. In an experiment, we use this compression strategy on three different transforms: the Fast Fourier transform, the Daubechies wavelet transform and the DOST. The results show that the DOST outperforms the two other methods.

Keywords: Image coding, spatial-frequency representation, DOST, multiresolution, Daubechies wavelets

1. INTRODUCTION

Image compression is an important step in many image-processing pipelines, allowing for smaller storage size, and faster download. Currently, JPEG image compression is one of the most pervasive image compression standards.¹ The most recent JPEG standard, called JPEG2000,² uses wavelets. Currently, wavelets are generally regarded as the leading technology for image compression.

Before wavelets, the Fourier transform (FT) was commonly used in image compression. The FT decomposes the image into its component frequencies, but does so globally so that each pixel affects every Fourier coefficient. Wavelets give a multiresolution decomposition in the spatial-scale domain. Even though the scale information can be approximately treated as frequency information (i.e. the fine scale information corresponds to the high frequency information and vice versa), the wavelet basis functions (e.g. the compactly supported Daubechies wavelets) are not entirely smooth. Hence, wavelet compression can be suboptimal on smooth parts of an image.

The Stockwell transform (ST) provides a continuous and infinitely differentiable kernel function and a full decomposition over the spatial-frequency domain. The orthonormal version of the Stockwell transform is the Discrete Orthonormal Stockwell transform (DOST), which gives a spatial-frequency decomposition with no redundancy. In this paper, we demonstrate the advantages of the DOST by analyzing the peak signal to noise ratio (PSNR) in an image compression experiment. We will see that a better approximation is achieved in the smooth areas of the image without sacrificing crisp edges.

Further author information: (Send correspondence to Yanwei Wang)

Yanwei Wang: E-mail: y72wang@math.uwaterloo.ca

Jeff Orchard: E-mail: jorchar@cs.uwaterloo.ca

2. REVIEW OF THE STOCKWELL TRANSFORM AND THE DOST

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

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

where f is the frequency, and t and τ are both time. In this way, the ST decomposes a signal into spatial (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 decomposition over frequency yields the Fourier coefficient of the signal, which highlights a special feature of the Stockwell transform. Hence, the original function $h(t)$ can be recovered by calculating the inverse Fourier transform of $H(f)$.

Using the equivalent frequency domain definition of the Stockwell transform, the discrete Stockwell transform (DST)³ can be written

$$S[jT, \frac{n}{NT}] = \sum_{m=0}^{N-1} H\left[\frac{m+n}{NT}\right] e^{-2\pi^2 m^2 / n^2} e^{i2\pi m j / N}, \quad (3)$$

for $n \neq 0$, and for the $n = 0$ voice,

$$S[jT, 0] = \frac{1}{N} \sum_{m=0}^{N-1} h[mT], \quad (4)$$

which is analogous to the DC value of the FT. The Stockwell transform has been used widely in various fields. For example, in geography it is used for analyzing internal atmospheric wave packets⁷, atmospheric studies⁸, characterization of seismic signals and global sea surface temperature analysis⁵. It is used in electrical engineering⁹, mechanical engineering¹⁰ and digital signal processing¹¹. It has also been applied in the medical field for human brain mapping¹², cardiovascular studies¹³, MRI analysis¹⁴ and the physiological effects of drugs¹⁵.

From (1), it is obvious that the Stockwell transform is an overcomplete representation. For a signal of length N , there can be N^2 Stockwell coefficients, but each one takes $O(N)$ to compute, suggesting that computing all N^2 coefficients of the ST has computational complexity of $O(N^3)$. That makes the Stockwell transform expensive to be applied to higher-dimensional cases. A more efficient basis is needed to pursue this spatial-frequency decomposition.

The DOST is a pared-down version of the fully redundant ST.¹⁶ 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, spans of high frequencies have higher sample rates.

Explicitly, the DOST basis vectors are defined as

$$S(kT)_{[\nu, \beta, \tau]} = \frac{1}{\sqrt{\beta}} \sum_{f=\nu-\beta/2}^{\nu+\beta/2-1} \exp\left(i2\pi \frac{k}{N} f\right) \exp\left(-i2\pi \frac{\tau}{\beta} f\right) \exp(-i\pi\tau), \quad (5)$$

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

$$S(kT)_{[\nu, \beta, \tau]} = ie^{-i\pi\tau} \frac{e^{-i2\alpha(\nu-\beta/2-1/2)} - e^{-i2\alpha(\nu+\beta/2-1/2)}}{2\sqrt{\beta} \sin \alpha}, \quad (6)$$

where $\alpha = \pi(k/N - \tau/\beta)$ is the center of the temporal window and ν indicates the center of each frequency band voice with a bandwidth of β .

To make the family of basis vectors in (6) orthonormal, the parameters, β , ν and τ , have to be picked suitably. Let the variable p represent the level of spatial decomposition. In Ref. 16, Dr. Stockwell defines DOST basis functions for each level p as

- $p = 0$, $S(kT)_{[\nu, \beta, \tau]} = 1$,
- $p = 1$, $S(kT)_{[\nu, \beta, \tau]} = \exp(-i2k\pi/N)$,
- $p = 2, \dots, \log_2(N) - 1$, pick:
 $\nu = 2^{(p-1)} + 2^{(p-2)}$,
 $\beta = 2^{(p-1)}$,
 $\tau = 0, \dots, \beta - 1$.

Mathematically, we can prove that these basis vectors are orthonormal,

$$\frac{1}{N} \int_0^N S_{[\nu', \beta', \tau']}[kT] S_{[\nu, \beta, \tau]}^*[kT] dk = \delta_{\nu' \nu} \delta_{\beta' \beta} \delta_{\tau' \tau}, \quad (7)$$

where

$$\delta_{xy} = \begin{cases} 1 & \text{for } x = y \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

is the Kronecker delta.

The advantage of the DOST over the DST is that it can represent a signal of length N with N coefficients.

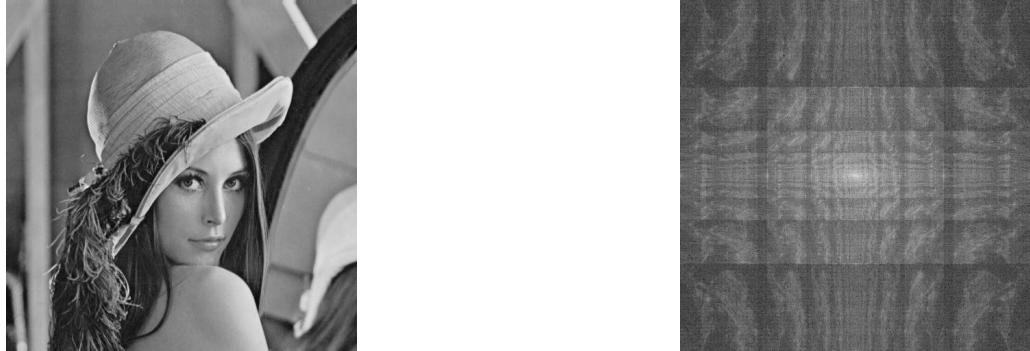


Figure 1. Lena and the logarithm of its DOST coefficients.

Figure 1 shows the logarithm of the DOST coefficients of one of our example images, Lena. As we can see, the coefficients decay very quickly, which makes the DOST a powerful tool for image compression.

3. METHODS

Our goal is to introduce the ST as a candidate tool for image compression. As an initial stab at the determining of the ST's capabilities, we compare it to two other transforms in a rudimentary compression methodology – simply dropping a percentage of the smallest coefficients (in modulus) and then reconstructing the images.

In our experiments, we used one of the most efficient families of wavelets, the Daubechies wavelets.¹⁷ The Daubechies wavelets form a family of orthogonal wavelets with a small number of coefficients. Generally, the order- K Daubechies wavelet has $2K$ non-zero coefficients, which makes the Daubechies wavelets efficient for image compression.¹⁸



Figure 2. Original sample images.

Table 1. PSNR for compression using 80% of coefficients.

Transform	Barbara	Lena	CT
DOST	90.27	88.87	86.69
FT	52.39	56.61	53.87
Haar	74.42	76.56	78.10
Daubechies-2	87.13	83.18	82.36
Daubechies-5	84.68	82.29	81.85
Daubechies-15	88.1665	84.59	80.22
Daubechies-38	81.9181	80.84	79.95

To compare the capabilities of the compression methods (DOST, FT, Daubechies), we conducted an experiment in which we applied each of the methods to three different test images (shown in Figure 2) at different compression rates. The test images are all 512×512 pixels in size.

Tables 1-3 report the PSNR of the compressed images for our experiment. In all cases, the DOST method yields a substantially higher PSNR and better quality than the FT and Daubechies methods. In addition (though not reported in the tables), the maximum intensity errors are roughly the same between the DOST and the Daubechies methods.

Figure 3 compares the original Barbara and different compressed versions using DOST, FT and Daubechies-2. As we can see, the DOST version remains sharper and keeps more detailed information (e.g. shadows behind the door, expression on the face and texture over the pants etc.) than the wavelet version.

Figures 4 and 5 show the intensity errors for the compression methods after compressing the Barbara images by 90% (reconstructing using only 10% of the coefficients). The distribution of the nonzero elements hints at each method's strengths and weaknesses. In particular, the FT method exhibits its largest errors in the regions containing high-frequency content. The DOST method shows relatively small errors throughout. Similar observations are made over different compression rates and different images.

4. CONCLUSION

According to our experiment, we conclude that, over the base line comparison, the DOST is a valuable tool for image compression by giving a higher PSNR and better quality than the wavelets and the FT. The distribution of the errors suggests that the DOST outperforms the wavelet compression in the smooth regions of the image. We understand that the state-of-the-art compression methods that use the FT and Wavelet method are considerably more complex than the simplistic compression method used in our experiments. This project is simply a pilot study, and more investigation is needed to assess the degree to which the DOST can challenge the best compression methods.

In the meantime, we are interested in investigating other uses of the DOST in image processing.

Table 2. PSNR for compression using 50% of coefficients.

Transform	Barbara	Lena	CT
DOST	55.17	53.20	52.45
FT	39.80	38.25	40.76
Haar	48.35	47.41	48.34
Daubechies-2	50.21	51.24	48.10
Daubechies-5	51.00	48.06	48.95
Daubechies-15	50.68	48.68	47.68
Daubechies-38	48.70	47.34	46.90

Table 3. PSNR for compression using 10% of coefficients.

Transform	Barbara	Lena	CT
DOST	34.31	33.40	33.25
FT	27.80	26.25	26.76
Haar	31.07	30.41	28.34
Daubechies-2	31.27	31.24	30.10
Daubechies-5	32.56	30.06	30.95
Daubechies-15	32.44	29.96	31.68
Daubechies-38	31.84	29.34	30.90

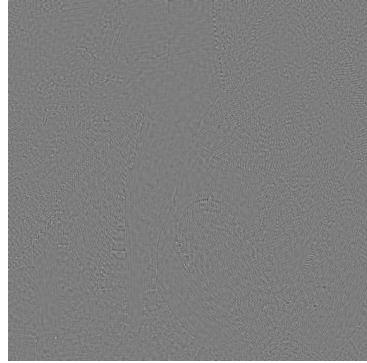
REFERENCES

- Wallace, G. K., "The JPEG still picture compression standard," *Source Communications of the ACM* **34**, 30–44 (April 1991).
- Christopoulos, C., Skodras, A., and Ebrahimi, T., "The JPEG2000 still image coding system: An overview," *Consumer Electronics, IEEE Transactions* **46**, 1103–1127 (November 2000).
- Stockwell, R. G., Mansinha, L., and Lowe, R. P., "Localization of complex spectrum: the S transform," *IEEE Transactions on Signal Processing* **144**, 998–1001 (April 1996).
- Mansinha, L., Stockwell, R. G., and Lowe, R. P., "Pattern analysis with two-dimensional spectral localisation: Application of two-dimensional S transforms," *Physica A* **239**, 286–295 (1997).
- Mansinha, L., Stockwell, R. G., Lowe, R., Eramian, M., and Schincariol, R. A., "Local S-spectrum analysis of 1-D and 2-D data," *Phys. Earth Planet. Inter.* **103**, 329C336 (November 1997).
- Eramian, M., Schincariol, R. A., Stockwell, R. G., Lowe, R. P., and Mansinha, L., "Review of applications of 1D and 2D S-transforms," *Proc. SPIE* **3078**, 558–568 (April 1997).
- Stockwell, R. G. and Lowe, R., "Aiglow imaging of gravity waves—1. Results from a small network of OH nightglow scanning imagers," *Journal of the Geophysical Research* **106**(D15), 185–203 (2001).
- Merzlyakov, E. G., Portnyagin, Y. I., Jacobi, C., Mitchell, N. J., Muller, H. G., Manson, A. H., Fachrutdinova, A. N., Singer, W., and Hoffmann, P., "On the longitudinal structure of the transient day-to-day variation of the semidiurnal tide in the mid-latitude lower thermosphere—i. Winter season," *Ann. Geophys.* **19**, 545–562 (2001).
- Dash, P., Panigrahi, B., and Panda, G., "Power quality analysis using S-transform," *IEEE Transactions on Power Delivery* **18** (April 2003).
- McFadden, P., Cook, J., and Forster, L., "Decomposition of gear vibration signals by the generalised S-transform," *Mechanical Systems and Signal Processing* **13**, 691–707 (September 1999).
- Pinnegar, C. R. and Mansinha, L., "The Bi-Gaussian S-transform," *SIAM Journal on Scientific Computing* **24**(5), 1678–1692 (2002).
- Andino, S., Menendez, R., Lantz, C., Blank, O., Michel, C., and Landis, T., "Non-stationary distributed source approximation: An alternative to improve localization procedures," *Human Brain Mapping* **14**, 81–95 (July 2001).

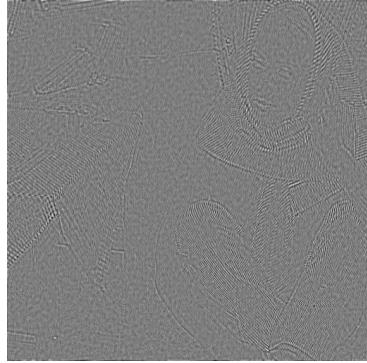


Figure 3. Original and compressed versions of Barbara using 10% of coefficients.

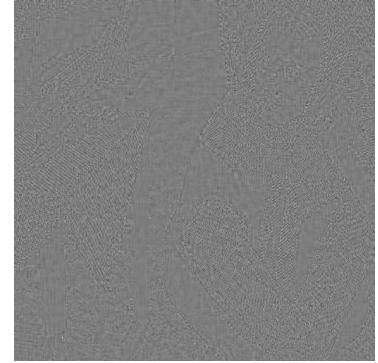
13. Varanini, M., Paolis, G., Emdin, M., Macerata, A., Pola, S., Cipriani, M., and Marchesi, C., “Spectral analysis of cardiovascular time series by the S-transform,” *Computers in Cardiology* , 383–386 (1997).
14. Goodyear, B. G., Zhu, H., Brown, R. A., and Mitchell, J. R., “Removal of phase artifacts from fMRI data using a Stockwell transform filter improves brain activity detection,” *Magnetic Resonance in Medicine* **51**(1), 16–21 (2004).
15. Assous, S., Humeau, A., Tartas, M., Abraham, P., and L’Huillier, J., “Physiological effects of indomethacin and celecoxib: An S-transform laser Doppler flowmetry signal analysis,” *Physics in Medicine and Biology* **50**, 1951–1959 (April 2005).
16. Stockwell, R. G., “A basis for efficient representation of the S-transform,” *Digital Signal Processing* **17**, 371–393 (January 2007).
17. Daubechies, I., “Orthonormal bases of compactly supported wavelets,” *Communications on Pure and Applied Mathematics*. **41**(7), 909–996 (1988).
18. Daubechies, I., [Ten Lectures on Wavelets], vol. 61, SIAM, University of Lowell, Philadelphia (1992).



(a) DOST

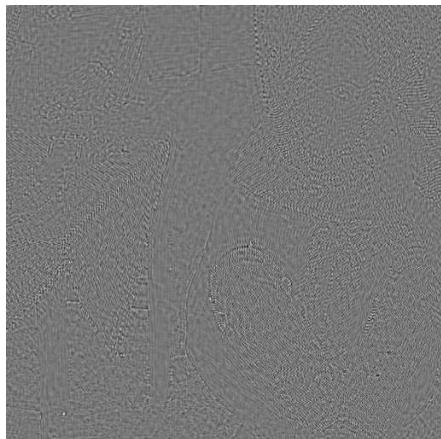


(b) FT

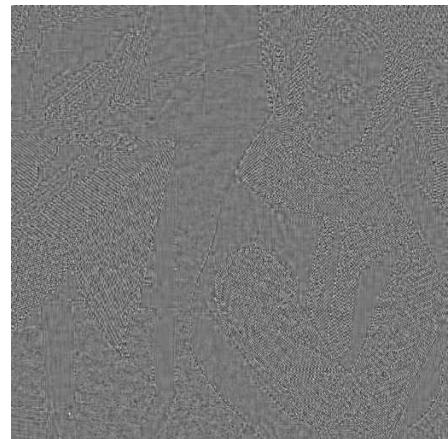


(c) Daubechies-2 wavelet

Figure 4. Intensity errors for Barbara image using 10% of coefficients (see Table 3). The gray level is set so that -80 maps to black and 82 maps to white (corresponding to the largest errors among these three methods).



(a) DOST



(b) Daubechies-2 wavelet

Figure 5. Intensity errors for Barbara image using 10% of coefficients (see Table 3). The gray level is set so that -45 maps to black and 51 maps to white (corresponding to the largest errors between the DOST and the Daubechies-2 wavelet method).