Available online at www.sciencedirect.com
SCIENCE DIRECT
ELSEVIER
Digital Signal Processing 15 (2005) 604–620
Digital Signal Processing
www.elsevier.com/locate/dsp

Time–frequency and time–time filtering with the S-transform and TT-transform

C.R. Pinnegar

Calgary Scientific, Calgary, AB, Canada

Available online 14 March 2005

Abstract

The limitations of frequency-domain filtering methods have motivated the development of alternative techniques, in which a filter is applied to a time–frequency distribution instead of the Fourier spectrum. One such distribution is the S-transform, a modified short-time Fourier transform whose window scales with frequency, as in wavelets. Recently it has been shown that the S-transform's local spectra have time-domain equivalents. Since each of these is associated with a particular window position on the time axis, collectively they give a time–time distribution. This distribution, called the TT-transform, exhibits differential concentration of different frequency components, with higher frequencies being more strongly concentrated around the localization position than lower frequencies. This leads to the idea of filtering on the time–time plane, in addition to the time–frequency plane. Examples of time–frequency filtering and time–time filtering are presented.

Keywords: Time–frequency analysis; Time–time analysis; Gabor transform; Wavelet transform; S-transform

1. Introduction

Standard Fourier-domain filtering techniques are constrained to stationary passbands and reject bands that are fixed for the entire duration of the time series. However, many important time series are nonstationary, such as electrocardiograms and recordings of human speech. Hence there is a need for filters with time-varying passbands and reject bands. Several methods of defining and implementing filters of this type have been developed during the past few decades [1–7]. One very intuitive approach [4] is to design a filter that

E-mail address: pinnegar@ucalgary.ca.

1051-2004/$ – see front matter
doi:10.1016/j.dsp.2005.02.002

can be applied to an invertible time–frequency representation (TFR) of the signal. This is conceptually similar to Fourier-domain filtering, except that the filter needs to be two-dimensional to accommodate the extra time dimension. There are a limitless number of candidate TFR's [8,9]. In Ref. [4], both the Wigner distribution [3,5,8], and the short-time Fourier transform (STFT) [8,10] (or Gabor transform) are considered. The Wigner distribution provides excellent resolution of the time–frequency content of chirp signals, but its cross terms lead to problematical artifacts for multicomponent time series. As a result, the Wigner distribution is not well suited to many practical applications. The STFT is often a better choice because it has no cross terms, but has a disadvantage because it uses the same analyzing window at all frequencies. Since a fixed window cannot accommodate the different cycle scales of different frequencies, resolution of low and high frequencies on the STFT is poor. The latter problem can be addressed by modifying the STFT to allow its window to scale with frequency, as in wavelets [6,12]. This gives the S-transform [11], whose general time-domain expression is

$$S[\tau, f] = \sum_{t=0}^{N-1} h[t] w[\tau - t, f] \exp\left(\frac{-2\pi i f t}{N}\right). \tag{1}$$

In (1), h is a discrete function of time, t and f are integer time and frequency indices (tT gives time in seconds, and f/NT gives frequency in Hz, where T is the sampling interval). The position of w (the S-transform window) on the t-axis is denoted by τ. Normally, w is defined so that its width scales inversely with frequency. One commonly used window is a Gaussian [11], denoted w_{GS},

$$w_{GS}[\tau - t, f] = \frac{|f|}{N\sqrt{2\pi}} \exp\left[\frac{-f^2(\tau - t)^2}{2N^2}\right]. \tag{2}$$

The width of the Gaussian part of w_{GS}, as measured between the peak and the point having $1/\sqrt{e}$ the peak amplitude, is equal to $N/|f|$, the wavelength of the fth Fourier sinusoid. Thus, at any f, w_{GS} always retains the same number of Fourier cycles. (In this paper, specific windows and their resulting S- and TT-transforms are denoted by suffices; the un-subscripted symbols w, S, and TT denote generalized forms of these quantities.) A sample S_{GS}, and its source time series, are shown in Figs. 1a and 1b.

An important constraint on w is the normalization condition

$$\sum_{\tau=0}^{N-1} w[\tau - t, f] = 1. \tag{3}$$

Due to (1) and (3), S collapses onto H, the discrete Fourier transform (DFT) of h, when summed over all possible values of τ,

$$H[f] = \sum_{\tau=0}^{N-1} S[\tau, f]. \tag{4}$$

The inverse DFT of (4) then gives the inverse S-transform,

$$h[t] = \frac{1}{N} \sum_{f=-N/2}^{N/2-1} \sum_{\tau=0}^{N-1} S[\tau, f] \exp\left(\frac{2\pi i f t}{N}\right). \tag{5}$$