

Article

Some Notes on the Use of the Windowed Fourier Transform for Spectral Analysis of Discretely Sampled Data

Robert W. Johnson

Alphawave Research, 29 Stanebrook Court, Jonesboro, GA 30238, USA;

E-Mail: robjohnson@alphawaveresearch.com; Tel./Fax: +1-770-477-0275

Received: 24 April 2013; in revised form: 20 May 2013 / Accepted: 21 May 2013 /

Published: 24 June 2013

Abstract: The properties of the Gabor and Morlet transforms are examined with respect to the Fourier analysis of discretely sampled data. Forward and inverse transform pairs based on a fixed window with uniform sampling of the frequency axis can satisfy numerically the energy and reconstruction theorems; however, transform pairs based on a variable window or nonuniform frequency sampling in general do not. Instead of selecting the shape of the window as some function of the central frequency, we propose constructing a single window with unit energy from an arbitrary set of windows that is applied over the entire frequency axis. By virtue of using a fixed window with uniform frequency sampling, such a transform satisfies the energy and reconstruction theorems. The shape of the window can be tailored to meet the requirements of the investigator in terms of time/frequency resolution. The algorithm extends naturally to the case of nonuniform signal sampling without modification beyond identification of the Nyquist interval.

Keywords: Fourier transform; Gabor transform; Morlet transform; multiresolution analysis

1. Introduction

The primary criticism leveled at the use of the continuous wavelet transform for the spectral analysis of discretely sampled data is that it fails to give quantitatively meaningful results. The power spectral density produced from the convolution of a wavelet basis and a discrete signal gives a qualitative picture of the temporal variation of its frequency content; however, when one attempts the reconstruction of the signal, the residual is not on the order of the precision of one's computational device. Likewise, the integrated margins of the power spectral density do not precisely equal the energy of the original signal. In other words, the discrete implementation of the continuous wavelet transform and its inverse do not

satisfy the energy and reconstruction theorems of spectral analysis. The goal of this investigation is to devise a multiresolution analysis that does satisfy those theorems.

Our insistence upon the satisfaction of the energy and reconstruction theorems is because they are first principle requirements related to the conservation of physical energy. When suitably defined, energy is not to be created nor destroyed; neither should it get lost in the shuffle. There is a deep relation between energy content and information content, as quantum mechanics teaches us, thus a loose grip on one implies a loose grip on the other. On a more practical level, the satisfaction of those theorems is among the requirements for a maximum entropy spectral analysis of data that includes the effects of measurement uncertainty, whose consideration is beyond the scope of this article.

The review by [1] remains a popular resource for practitioners of wavelet analysis. It relies on the method by [2] for the reconstruction of the data signal. [3] are credited with establishing the reconstruction theorem in the continuum. [4] and [5] developed the theory of multiresolution analysis, and [6] constructed the first orthonormal basis with compact support, leading to the implementation of the dyadic wavelet transform in terms of finite impulse response digital filter banks [7]. Some interesting applications of the continuous wavelet transform for the spectral analysis of data can be found in references [8–12].

This paper is organized as follows. First, we will quickly review the theory of the continuous Fourier transform and its discrete implementation. Next, we will look at the Gabor transform and its relation to the wavelet transform using the Morlet basis. We will then propose an algorithm for a layered window transform whose spectral density is similar to that of the Morlet transform yet which satisfies the energy and reconstruction theorems. Following that, we will demonstrate that the algorithm works unaltered for data with irregular sampling once the corresponding Nyquist interval is identified. Finally, we will look at how the selection of the window affects the time/frequency resolution of the transform. We will conclude with a brief summary and suggestions for applications.

Some notations used throughout this paper are explained here. Scalars are written as s , while vectors and matrices are written as \mathbf{v} and \mathbf{M} respectively and may be defined in terms of their components, e.g., $\mathbf{v} \equiv v_s = v(s)$. The transpose of a vector or matrix is indicated by the superscript \mathbf{v}^T , and the conjugate transpose by \mathbf{M}^\dagger , using the standard rules for matrix multiplication. Inner products may be written in bra-ket notation as $\langle a|b\rangle_s$, and matrix entries as the ket-bra $|a\rangle\langle b|$. Expectation values are notated as $\langle v(s)\rangle_s$. Sets will be indicated by their boundaries $[a, b]$, and whether continuous or discrete must be derived from context. The operation of rounding up, *i.e.*, taking the next greatest integer, is denoted by $\lceil a \rceil$, and when the notation $a \rightarrow b$ is encountered, it is understood to mean “ a is replaced by b ”. The programs used for this analysis are available online [13].

2. Continuous and Discrete Fourier Transforms

Let us begin by considering the Fourier transform in the continuum over axes of time and frequency. Suppose we have some signal $y(t)$, possibly complex valued, of finite energy $E_y \equiv \langle y|y\rangle_t = \int_{-\infty}^{\infty} y^*(t)y(t)dt < \infty$. If the signal carries units of u_y , then the signal energy has units of $u_E \equiv u_y^2 u_t$. The units of signal energy are proportional to those of physical energy in joules by a factor of the load impedance $E_J = E_y/Z_L$; for example, if the signal has units of volts $u_y = V$ and time is measured

in seconds $u_t = s$ so that $u_E = V^2s$, then accounting for the load impedance in ohms $u_Z = \Omega$ gives physical units of $V^2s/\Omega = (ml^2/t^2q)^2t/(ml^2/tq^2)$ in terms of mass, length, time, and charge, which is equivalent to joules. The units of frequency are reciprocal to those of time $u_f = u_t^{-1}$, and the term “power” is understood to mean “energy density” and must be qualified by the domain for its distribution. The squared modulus $|y(t)|^2 \equiv P(t)$ can thus be identified as the temporal power of the signal.

Under suitable conditions not elaborated here, the Fourier integral and its inverse,

$$\hat{y}(f) \equiv \int_{-\infty}^{\infty} \exp(-i2\pi ft)y(t)dt, \quad \hat{y}(t) \equiv \int_{-\infty}^{\infty} \exp(i2\pi ft)\hat{y}(f)df \tag{1}$$

define an integral transform pair satisfying the Plancherel energy theorem $\int_{-\infty}^{\infty} |y(t)|^2 dt = \int_{-\infty}^{\infty} |\hat{y}(f)|^2 df$ and the Fourier inversion theorem $\hat{\hat{y}}(t) = y(t)$. The representation over the frequency axis $\hat{y}(f)$ carries units of $u_{\hat{y}} = u_y u_t$ so that the squared modulus $|\hat{y}(f)|^2 \equiv P(f)$ gives the spectral power of the signal, and the spectral energy is $E_{\hat{y}} \equiv \langle \hat{y} | \hat{y} \rangle_f = \int_{-\infty}^{\infty} \hat{y}^*(f)\hat{y}(f)df$. When the signal is of finite duration $t \in [0, T]$, the Fourier transform’s response to a component sinusoid $y(t) = \exp(i2\pi f't)$ is

$$\hat{y}(f) = \int_0^T \exp[-i2\pi(f - f')t]dt = \exp(-i\pi f_{\Delta}T) \frac{\sin(\pi f_{\Delta}T)}{\pi f_{\Delta}} \tag{2}$$

expressed in terms of the frequency offset $f_{\Delta} \equiv f - f'$, thus the spectral power of the uniform window function $\Phi_T(t) = T^{-1/2}$ normalized to unit energy describes the leakage in the frequency spectrum induced by the finite duration of the signal,

$$L_T(f_{\Delta})df_{\Delta} \equiv |\hat{\Phi}_T(f_{\Delta})|^2 df_{\Delta} = T^{-1} \left[\frac{\sin(\pi f_{\Delta}T)}{\pi f_{\Delta}} \right]^2 df_{\Delta} \tag{3}$$

with domain $f_{\Delta} \in [-\infty, \infty]$. The continuum leakage function is not periodic in f_{Δ} , as its magnitude decays according to f_{Δ}^{-2} .

With regard to practical data analysis, let us suppose now that the signal is given in terms of discrete samples in time with uniform duration and spacing. (Those requirements will be relaxed in Section ??.) The time axis can then be described in terms of integers $t \in [1, T]$ with unit $u_t \equiv \Delta_t$ given by the sample rate, allowing the signal to be written as a vector $\mathbf{y} \equiv y_t = y(t)$. The signal energy can be expressed in terms of matrix multiplication as $E_y = \mathbf{y}^\dagger \mathbf{D}_t \mathbf{y}$, where the temporal metric $\mathbf{D}_t = I_T \Delta_t$ is proportional to the identity matrix of order T . The effect of frequency aliasing induced by the uniform sampling is often misunderstood. If we evaluate the Fourier transform of the discrete window with unit energy,

$$\hat{\Phi}_T(f_{\Delta}) = T^{-1/2} \sum_{t=1}^T \exp(-i2\pi f_{\Delta}t)\Delta_t = T^{-1/2} \exp[-i\pi f_{\Delta}(T + 1)] \frac{\sin(\pi f_{\Delta}T)}{\sin(\pi f_{\Delta})} \Delta_t \tag{4}$$

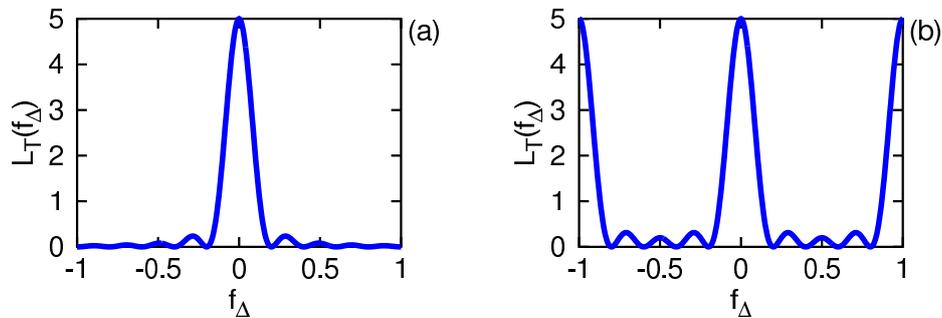
we find that the leakage function becomes periodic in f_{Δ} ,

$$L_T(f_{\Delta})df_{\Delta} \rightarrow T^{-1} \left[\frac{\sin(\pi f_{\Delta}T)}{\sin(\pi f_{\Delta})} \right]^2 \Delta_t^2 df_{\Delta} \tag{5}$$

with period 1 in units of $u_f = \Delta_t^{-1}$. The principle branch is usually chosen as $f \in [-1/2, 1/2]\Delta_t^{-1}$, but one should always respect the periodic nature of the spectrum. The normalization of the spectral power likewise is defined over one period of the frequency axis, $\sum_{t=1}^T |y(t)|^2 \Delta_t = \int_{-1/2}^{1/2} |\hat{y}(f)|^2 df$. In

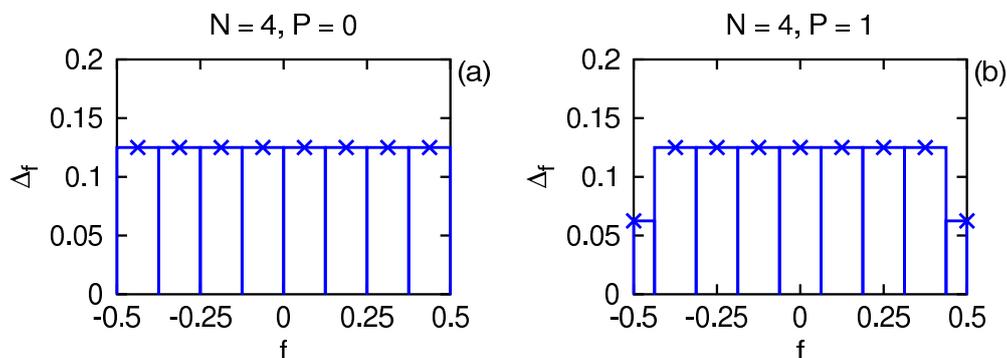
Figure ?? we compare the continuum and discrete leakage functions for a window with duration $T = 5$, where we see the expected oscillatory lobes as well as the periodicity induced by the discrete sampling in time.

Figure 1. Comparison of the continuous time leakage function in (a) to its discrete time counterpart in (b) for a uniform window of duration $T = 5$ over a continuous frequency axis.



Now let us consider the discretization of the frequency axis into uniform bins over its principle branch $f \in [-f_c, f_c]$, where $f_c \equiv (2\Delta_t)^{-1}$ is the Nyquist critical frequency. To fully specify the frequency metric D_f , one must state its order N , equal to the number of positive frequencies $\leq f_c$, as well as its parity $P \in \{0, 1\}$, indicating whether an even or odd number of bins $N' = 2N + P$ is used to span the domain. When $P = 1$, the bins on the edges corresponding to frequencies $\pm f_c$ have a bin width equal to one half that of the others, so that the limits of integration are respected; because of aliasing, the integrand will have the same value at those locations, so that only one full bin's contribution is counted. This convention differs from that usually given for the discrete Fourier transform [?], where the factor of 1/2 is absorbed by the coefficients rather than the metric. The order N specifies the spacing of the frequencies as $\Delta_f = (2N\Delta_t)^{-1}$, and the metric can be written as $D_f/\Delta_f = I_{N'} - (|1\rangle\langle 1| + |N'\rangle\langle N'|)P/2$ so that $\text{Tr } D_f = \Delta_t^{-1}$. The central frequencies of the bins can be expressed as $f_n = [n - (N' + 1)/2]\Delta_f$ for integers $n \in [1, N']$. Figure ?? compares the even and odd discretizations of order $N = 4$. While the odd parity discretization is perhaps more familiar, for many of our purposes the even discretization will prove more convenient.

Figure 2. Comparison of the even (a) and odd (b) frequency discretizations of order $N = 4$ with the central frequencies indicated by \times .



We are finally ready to define the two-sided discrete Fourier transform (of order N and parity P) of a discretely sampled signal (of duration T). If we collect our Fourier basis functions into the form of

a matrix, $\Theta \equiv \Theta(t, n) = \exp(i2\pi f_n t)$, then we can easily write the discrete transform pair in terms of matrix multiplication, $\hat{\mathbf{y}} \equiv \Theta^\dagger \mathbf{D}_t \mathbf{y}$ and $\hat{\hat{\mathbf{y}}} \equiv \Theta \mathbf{D}_f \hat{\mathbf{y}}$, as well as the spectral energy $E_{\hat{\mathbf{y}}} = \hat{\mathbf{y}}^\dagger \mathbf{D}_f \hat{\mathbf{y}}$. In component notation, one has

$$\hat{y}(n) \equiv \sum_{t=1}^T \exp(-i2\pi f_n t) y(t) \Delta_t, \quad \hat{\hat{y}}(t) \equiv \sum_{n=1}^{N'} \exp(i2\pi f_n t) \hat{y}(n) \Delta_f \tag{6}$$

where Δ_f is understood to account for the edge bins if $P = 1$. If the signal $y(t)$ is real, so that $\hat{y}(-f) = \hat{y}(f)^*$, then one may define the one-sided transform, retaining only the nonnegative portion of the frequency axis with $N' \rightarrow N + P$ bins and $f_n \rightarrow [n - (P + 1)/2] \Delta_f$, by renormalizing the basis functions to twice the energy via $\Theta \rightarrow \sqrt{2} \Theta$ and letting $\hat{\hat{y}}(t) \rightarrow \text{Re } \hat{\hat{y}}(t)$. Hereafter, we will assume $y(t)$ is real so that we can focus our attention on the one-sided spectrum.

In the fully discrete setting, the satisfaction of the energy and reconstruction theorems is dependent upon there being a sufficient number of degrees of freedom (DOF) in the basis functions to fully represent the information content of the signal. For real $y(t)$, the number of DOF is equal to the signal duration T . For the one-sided transform, whether of even or odd parity, the number of DOF is equal to $2N$, as each frequency's coefficient has both an amplitude and a phase, $\hat{y}_n = A_n \exp(i\omega_n)$, except for the cases $f_n = 0$ or f_c that have only an amplitude. The critical frequency f_c is identified as the lowest positive frequency whose basis function is entirely real over the discrete time axis, an observation that will prove useful when we consider the case of irregularly sampled data. The minimal order beyond which the energy and reconstruction theorems are satisfied can be evaluated as $N_{\min} = \lceil T/2 \rceil$ for a real signal (and $N_{\min} = T$ for a complex one). That condition is realized (for even T) when $\Theta^\dagger \Theta \Delta_t = \mathbf{I}_{N'} T$, indicating that the basis functions Θ form an orthogonal set for $N = T/2$. However, orthogonality over the discrete time axis is not a requirement, as $\mathbf{y}^\dagger \mathbf{D}_t \mathbf{y} = \hat{\mathbf{y}}^\dagger \mathbf{D}_f \hat{\mathbf{y}}$ and $\mathbf{y} = \hat{\hat{\mathbf{y}}}$ exactly (meaning to the precision of one's computational device) for any $N \geq N_{\min}$ and either parity P .

Lastly, let us look at what happens when one tries to express the fully discrete, one-sided Fourier transform over an axis of period (or scale) s rather than frequency f . In the continuum, the relation between the axes $f = s^{-1}$ yields the Jacobian $|df/ds| = s^{-2}$, thus the spectral power over scale $P(s) \equiv |\hat{y}(s)|^2$ is equal to the spectral power over frequency $P(f)$ multiplied by the Jacobian, $P(s) = f^2 P(f)$, otherwise expressed as $\hat{y}(s) = f \hat{y}(f)$. In the discrete setting, however, one must be explicit with the mapping of the bin boundaries so that $\sum_n |\hat{y}(f_n)|^2 \Delta_f = \sum_n |\hat{y}(s_n)|^2 \Delta_s$. Let $\Delta_f = b - a$ be the width of one frequency bin centered on $f_n = (a + b)/2$, which gets mapped to the scale bin $\Delta_s = a^{-1} - b^{-1}$ such that $\Delta_f/\Delta_s = ab$ and $P(s_n) = abP(f_n)$. One may very well ask, in what sense is $s_n = f_n^{-1}$ the center of the scale bin? The value f_n is both the mean and the median of the frequency bin with uniform measure $p^f = \Delta_f^{-1}$, but the mean period over that bin is $\langle f^{-1} \rangle_f = \log(b/a) \Delta_f^{-1}$. The value s_n is recognized as the median of the scale bin with measure $p^s = s^{-2} \Delta_f^{-1}$, such that $\int_{1/a}^{2/(a+b)} p^s ds = 1/2$. The edges for odd parity are handled similarly with respect to the half-width bins. For either parity, the bin whose lower boundary is $a = 0$ is the one that causes trouble, as $\Delta_s = \infty$ in that case. One can contrive to neglect that bin for odd parity by subtracting the mean of the signal $y(t) \rightarrow y(t) - \langle y(t) \rangle_t$, but that procedure does not repair the problem for even parity.

To illustrate the difficulty, in Figure ?? we display the mapping of the spectral power for a signal with duration $T = 20$, a sum of two sinusoids at frequencies 1/40 and 1/4 normalized to unit energy $E_y = 1$.

In panels (a) and (b) we show the mapping for the one-sided transform of order $N = 20$ and odd parity over domain $f \in [0, f_c]$. While the higher frequency's contribution remains apparent, that of the lower frequency has been washed away by the measure factor. The DOF carried by the lowest frequency bin are unrecoverable from the mapping over scale on account of the infinite bin width. In panels (c) and (d) we repeat the procedure but for domain $f \in [f_c, 2f_c]$. This time, all the DOF are properly mapped, thus the energy and reconstruction theorems are satisfied. In Table ?? we give some numerical results from the evaluation shown in the figure. In short, there is nothing to be gained by working on the scale axis rather than frequency other than a headache in dealing with the effect of aliasing, as one requires the same set of basis functions Θ and minimal order N_{\min} to satisfy the fundamental theorems of spectral analysis in either case.

Figure 3. Comparison of the mapping over odd parity axes of frequency f and scale $s = f^{-1}$ of the discrete spectral power of a signal with two sinusoidal components for $f \in [0, f_c]$ in (a) and (b) and for $f \in [f_c, 2f_c]$ in (c) and (d) with the central frequencies indicated by \times .

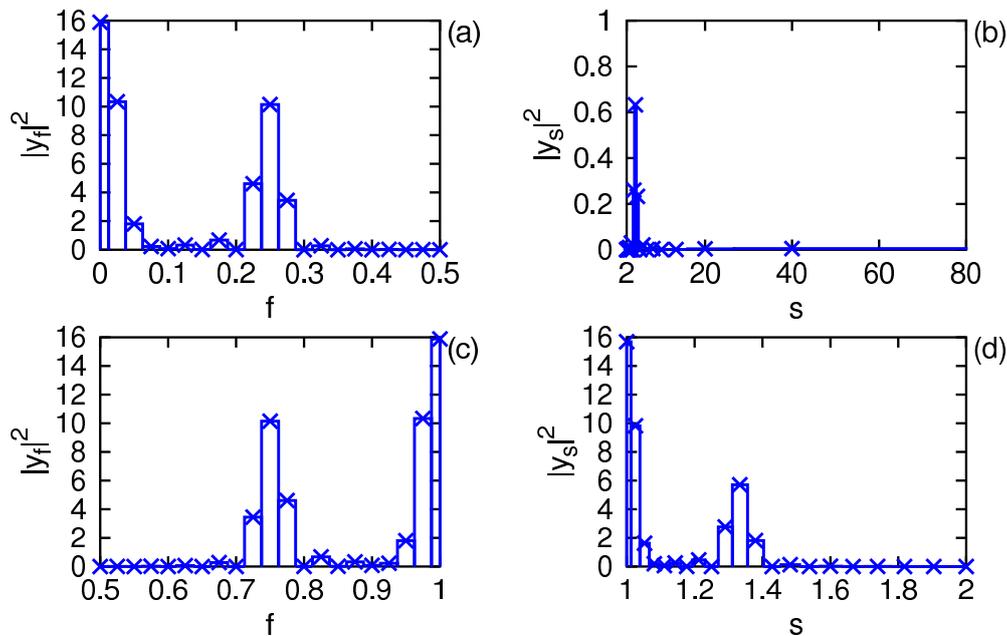


Table 1. Numerical results from the evaluation of Figure ??.

P	$f \in [0, f_c]$				$f \in [f_c, 2f_c]$			
	$\text{Tr } \mathbf{D}_f$	$\sum_n P(f_n)\Delta_f$	$\text{Tr } \mathbf{D}_s$	$\sum_n P(s_n)\Delta_s$	$\text{Tr } \mathbf{D}_f$	$\sum_n P(f_n)\Delta_f$	$\text{Tr } \mathbf{D}_s$	$\sum_n P(s_n)\Delta_s$
0	0.5	1.0	INF	NAN	0.5	1.0	1.0	1.0
1	0.5	1.0	INF	NAN	0.5	1.0	1.0	1.0

3. Gabor and Morlet Transforms

Let us turn now to the consideration of how to transform the signal $\mathbf{y} \equiv y(t)$ into a time-frequency representation $\hat{Y} \equiv \hat{Y}(f, t)$ by modulating the exponential oscillations of the Fourier basis with a discrete window function, $\Theta(f, t) \rightarrow \Psi(f, t) \equiv \Phi(t)\Theta(f, t)$. The Gabor transform is defined by the

use of a Gaussian window with decay parameter σ , which we will notate as $\Phi(t) \propto \exp^{-\pi/2}(t^2/\sigma^2) \equiv e^{-\pi t^2/2\sigma^2}$. Its normalization to unit energy $\Phi \rightarrow \Phi/\langle\Phi|\Phi\rangle_t^{1/2}$ depends upon the duration of the window, parametrized by its half-width τ . In the continuum, one can write $\int_{-\infty}^{\infty} \Phi^2(t)dt \propto \sigma$ and $\int_{-\tau}^{\tau} \Phi^2(t)dt \propto \sigma \operatorname{erf}(\pi^{1/2}\tau/\sigma)$; however, in the discrete setting with index $t \in [-\tau, \tau]$, the window energy $\sum_t \Phi^2(t)\Delta_t$ is not expressible in closed form. For the one-sided spectrum, $\Phi \rightarrow \sqrt{2}\Phi$ so that its energy equals 2. To fully specify which Gabor transform is being used, one must state the values of N , P , σ , and τ . We will return to the question of finding the minimal order of the Gabor transform for a signal of duration T in Section ??; for now, we will assume that N is greater than N_{\min} of the previous section so that the fundamental theorems are satisfied.

There are two alternatives for the definition of the phase convention in the transform, which differ in whether the phase is expressed relative to the origin of the signal’s time axis or that of the window:

$$\hat{Y}(n, \hat{t}) \equiv \sum_{t'=-\tau}^{\tau} \Phi(-t') \exp^{i2\pi}[-f_n(\hat{t} + t')]y(\hat{t} + t')\Delta_t \tag{7}$$

$$\hat{y}(t) \equiv \sum_{n=1}^{N'} \exp^{i2\pi}(f_n t) \sum_{t'=-\tau}^{\tau} \Phi(t')\hat{Y}(n, t + t')\Delta_t\Delta_f \tag{8}$$

where the time axis for the transform coefficients carries an index $\hat{t} \in [1 - \tau, T + \tau]$, or else

$$\hat{Y}(n, \hat{t}) \equiv \sum_{t'=-\tau}^{\tau} \Phi(-t') \exp^{i2\pi}(-f_n t')y(\hat{t} + t')\Delta_t \tag{9}$$

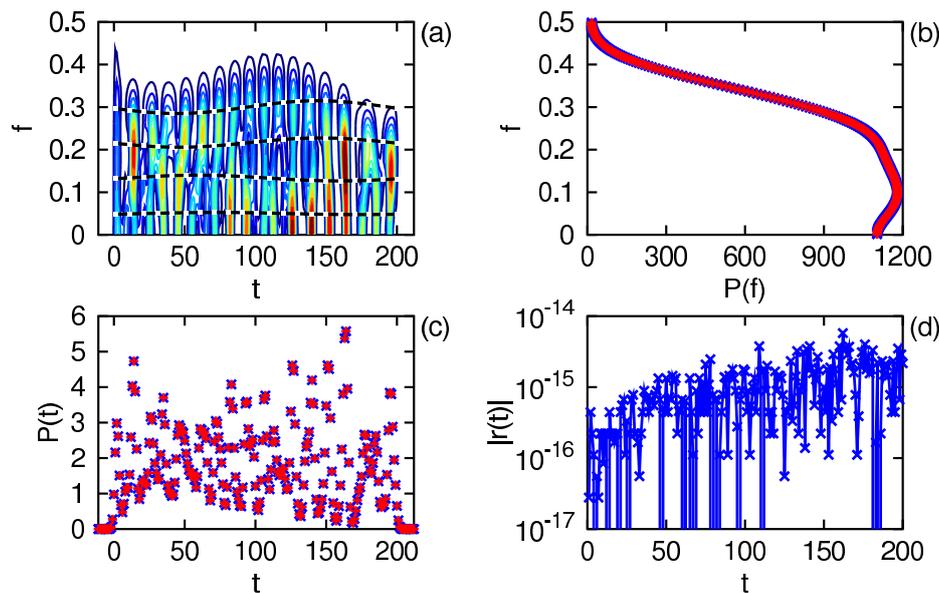
$$\hat{y}(t) \equiv \sum_{n=1}^{N'} \sum_{t'=-\tau}^{\tau} \Phi(t') \exp^{i2\pi}(-f_n t')\hat{Y}(n, t + t')\Delta_t\Delta_f \tag{10}$$

The latter is more familiar, but either is equally valid in terms of satisfying the energy and reconstruction theorems. The sign of the argument to the window function is chosen deliberately so that these expressions remain valid for the case of a window that is not symmetric in t' . The signal is zero padded for values of $\hat{t} + t'$ outside the original domain of $[1, T]$, as no other value could be assigned without changing the energy hence information content of the signal. For either phase convention, the temporal spectral power is defined as $P(n, \hat{t}) \equiv |\hat{Y}(n, \hat{t})|^2$ with units of energy per time per frequency. (The normalization of the window introduces a factor of $\Delta_t^{-1/2}$ so that \hat{Y} carries units of $u_y u_t^{1/2}$.) For comparison with the signal energy E_y , let us evaluate the spectrum’s energy as $E_{\hat{Y}} \equiv \sum_n \sum_{\hat{t}} P(n, \hat{t})\Delta_t\Delta_f$ and the reconstruction’s energy as $E_{\hat{y}} \equiv \sum_t |\hat{y}(t)|^2 \Delta_t$.

For this section and the next, let us consider a particular real signal $y(t)$ comprised of 4 sinusoids with unit amplitude and non-stationary frequency of duration $T = 200$. The variation in the instantaneous frequencies is assigned an amplitude of 5% and a period of T . Phases for the oscillations and variations are selected randomly. In Figure ?? we show the results of the one-sided Gabor transform with $N = 200$ and $P = 0$ using window parameters of $\sigma = 2\sqrt{\pi}$ and $\tau = 12$. Panel (a) displays the contours of the energy density $P(n, \hat{t})$ over axes of time and frequency. Also indicated are the instantaneous frequencies used to generate the data. Panels (b) and (c) display the marginal energy densities over axes of frequency and time, respectively. The marginal energy densities are evaluated by taking each sum over an axis independently, *i.e.*, $P(n) \equiv \sum_{\hat{t}} P(n, \hat{t})\Delta_t$ and $P(\hat{t}) \equiv \sum_n P(n, \hat{t})\Delta_f$. Panel (d) gives the absolute

value of the residual $r(t) \equiv \hat{y}(t) - y(t)$, which is on the order of the machine precision $\sim 10^{-16}$. The ratios $E_{\hat{y}}/E_y$ and $E_{\hat{y}}/E_{\hat{y}}$ are equal to unity to the same precision, thus the fundamental theorems are satisfied. At these parameter values, the spectral resolution is poor while the temporal resolution is sharp.

Figure 4. One-sided Gabor transform of a real signal as described in the text. The power spectrum is shown in (a) with the signal’s instantaneous frequencies indicated by the dashed lines. The marginal spectral power is shown in (b) as \times , as is the marginal temporal power in (c), and each is compared with the convolution of the window and signal energies indicated by $+$. The absolute value of the reconstruction residual is displayed in (d).



The marginal densities of the Gabor power spectrum can be compared with the convolution of the window’s energy density with that of the signal in either the temporal or spectral representations. Looking first at the temporal representation, one finds the relation

$$P(\hat{t}) = \sum_{t'=-\tau}^{\tau} \Phi^2(-t')y^2(\hat{t} + t')\Delta_t \tag{11}$$

where Φ in this case is normalized to unit energy. In the spectral representation, the convolution is taken in the modular (periodic) sense, most easily performed using the two-sided transform. Let $\hat{y}(f_n)$ be the two-sided Fourier transform of the signal with order N and parity P , and let $\hat{\Phi}(f'_m)$ be the transform of the window function, again normalized to unit energy, with the same order and odd parity such that $M = 2N + 1$. Then $P(n)$ in the two-sided sense $n \in [1, 2N + P]$ can be written

$$P(n) \approx \sum_{m=1}^M |\hat{\Phi}(M + 1 - m)|^2 |\hat{y}[\text{mod}(n + m - N - 2, 2N) + 1]|^2 \Delta_f \tag{12}$$

where $\text{mod}(a, b) \equiv a \text{ mod } b$ and with respect to the half-width of the edge bins, from which the values corresponding to the one-sided transform can be extracted and doubled. The relation above is written as an approximation because there is some subtlety to the evaluation of the RHS.

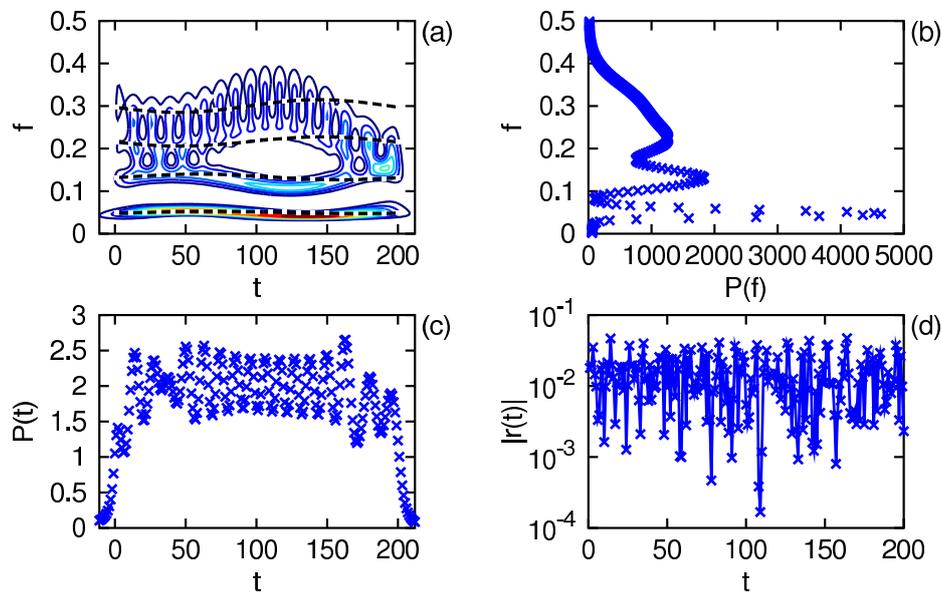
So far, we have imposed no condition on the window duration $T' = 2\tau + 1$ other than being odd for integer τ , and indeed, there is none. In the lower limit $\tau = 0$, the basis functions are constant,

$\Psi(0, n) = \sqrt{2}$ (or 1 for the two-sided transform), and all frequency resolution is lost, yet the fundamental theorems are still satisfied for any order $N \geq 1$. In the opposite extreme, for example $\tau = T$ such that $T' \gg T$ with $\sigma = \infty$, the basis functions become essentially those of the Fourier transform with minimal order N_{\min} . Simply put, the duration of the window in the short-time Fourier transform does not need to be short. When the temporal bandwidth of the window approaches or exceeds that of the signal, the evaluation of the RHS in Equation (??) becomes suspect if the order N is not sufficient to resolve both the signal and the window; the practical solution is to increase N . Furthermore, one can verify that the Gabor transform remains well behaved (if not particularly useful) for the case $\sigma \rightarrow i\sigma$ so that the window grows exponentially rather than decaying. The only requirement is that the window be real valued and normalized explicitly, *i.e.*, discretely, to unit energy, times 2 for the one-sided transform.

We can now introduce the Morlet transform by promoting the window parameters from constants to functions of the central frequency, $\Phi(t) \rightarrow \Phi_n(t)$. Let $\sigma_n \equiv \sigma/f_n$, and with respect to the discrete sampling in time, let $\tau_n \equiv \lceil \tau/f_n \rceil$ for constants σ and τ . The parameter τ itself does not need to be an integer, as τ_n is what defines the window duration for bin n . A difficulty with τ_n arises immediately for odd parity when $f_n = 0$; one can ameliorate the situation by using τ_n from the lowest positive bin in that case. Again, either the phase convention of Equations (?? and ??) or (?? and ??) can be used. In Figure ?? we show the results of the one-sided Morlet transform of order 200 and even parity with parameters $\sigma = \sqrt{\pi}$ and $\tau = 6$ chosen so that the window of the previous Gabor transform corresponds to the Morlet window at the critical frequency f_c . The power spectrum in panel (a) has the marginal densities shown in panels (b) and (c); however, this time the ratios $E_{\hat{Y}}/E_y$ and $E_{\hat{y}}/E_{\hat{Y}}$ differ from unity on the order of 1%, interestingly by nearly the same value. Likewise, the absolute value of the residual displayed in panel (d) is on the order of 1%, comparable to the Farge method [?] as reported by Torrence and Compo [?]. While the lowest frequency component has been well resolved, the spectral resolution of the upper frequencies remains poor. Furthermore, without a single window in operation, one cannot perform the comparison of the marginal densities with the convolution of the window and signal energy densities corresponding to Equations (??) and (??).

Let us insert here some comments on what is called the admissibility condition. The statement often is made by authors that the wavelet must have a mean of zero so that its Fourier coefficient at $f = 0$ vanishes. While that remark might apply in the continuum, it is not appropriate for discretely sampled wavelets of finite duration. If one examines in detail the proof of the admissibility condition [?], one finds that it relies crucially on the property of continuity. In full, the admissibility condition states that if the wavelet is *continuous* in time with *continuous* Fourier transform, then its mean must be zero. In the context considered here, neither of those conditions is met: In the temporal representation the wavelet is a piece-wise constant function with jump discontinuities at the edges of the temporal bins, and likewise for the frequency representation that, while of arbitrary resolution beyond the minimal order, must nonetheless be evaluated discretely for any practical analysis of data. If one computes the leakage functions for the discrete Morlet basis [?], one finds that an artificial discontinuity is induced by the imposition of the zero mean condition. As a final argument against subtracting the mean of the discrete basis functions, note that the Gabor transform works perfectly well for any values of the window parameters without including such procedure.

Figure 5. One-sided Morlet transform of a real signal as described in the text. The power spectrum is shown in (a) with the signal’s instantaneous frequencies indicated by the dashed lines. The marginal spectral power is shown in (b) as ×, as is the marginal temporal power in (c). The absolute value of the reconstruction residual is displayed in (d).



There are many suggestions found in the literature for improving the performance of the discretely implemented continuous wavelet transform. Let us examine a few of them here, with their absolute residuals displayed in Figure ?? and a numerical summary in Table ?. For panel (a) the mean of the signal is subtracted before entering the spectral analysis. As one can see, this procedure results in few changes, but for consistency of comparison the remaining panels will also use the mean subtracted signal. Noticing that the ratios $E_{\hat{\gamma}}/E_y$ and $E_{\hat{\delta}}/E_{\hat{\gamma}}$ are very nearly equal, one can renormalize the spectrum and reconstruction *a posteriori* by their geometric mean $C = (E_{\hat{\delta}}/E_y)^{1/2}$, yielding the residual shown in panel (b). This renormalization prescription works best when there is not much energy in the upper portion of the frequency axis [?]. Perhaps more familiar is the prescription to shift the central frequency $2\pi \rightarrow \omega_1$. At these window parameter values, a shift of $\omega_1/2\pi = 1.0127$ is found to work well [?], giving the residual shown in panel (c). If one then applies the renormalization prescription, one gets the residual displayed in panel (d). While there has been a noticeable improvement of three orders of magnitude, the root-mean-square residual remains far above the machine precision.

Figure 6. Absolute residuals for several methods of improving the transform response as described in the text. In panel (a), the mean of the signal is subtracted; in panel (b), the mean is subtracted with renormalization of the coefficients; in panel (c), the mean is subtracted with a shift of central frequency; and in panel (d), the mean is subtracted with frequency shift and renormalization.

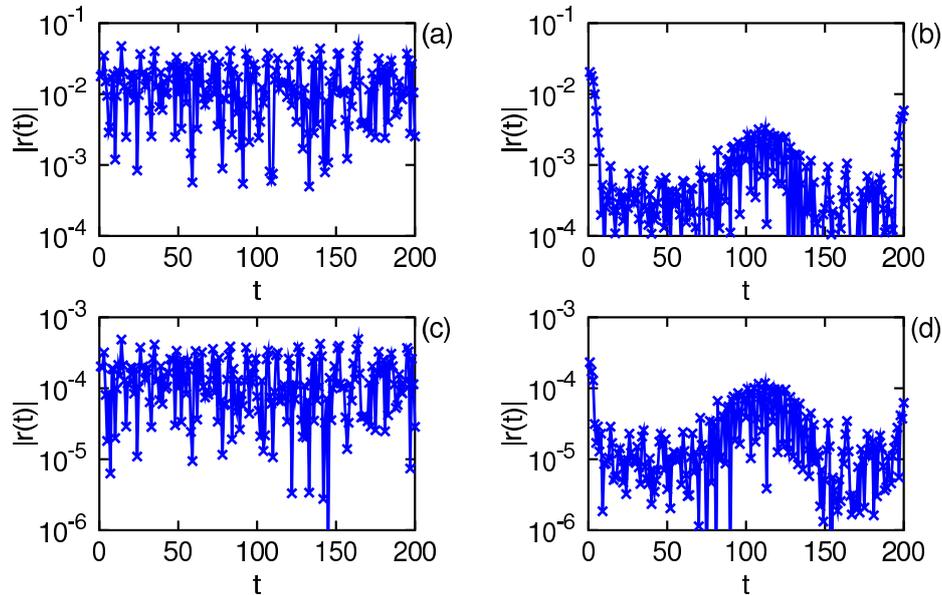


Table 2. Numerical results from the evaluation of Figure ??.

Panel	(a)	(b)	(c)	(d)
$\langle r^2(t) \rangle_t^{1/2}$	1.8215e -02	2.6142e -03	1.8213e -04	4.1799e -05
$E_{\hat{Y}}/E_y - 1$	1.2885e -02	-1.7468e -06	1.2674e -04	-4.4658e -10
$E_{\hat{y}}/E_{\hat{Y}} - 1$	1.2889e -02	1.7468e -06	1.2674e -04	4.4658e -10

4. Layered Window Transform

The primary feature of the Morlet transform is that it uses a different window function Φ_n for each bin along the frequency axis. That is also its main difficulty, as each bin corresponds to a particular instance of the Gabor transform, any of which would work fine independently, but when the separate bins are isolated and combined, there is no reason to suppose that their marginal sums will equal the signal energy. The correspondence between the Morlet and Gabor transforms becomes more apparent if one redefines the window decay σ_n relative to the window width τ_n rather than the central frequency f_n , in which case regions along the frequency axis of the Morlet spectrum are drawn from a particular Gabor spectrum. Simply pasting together pieces from different Gabor transforms from the outset appears doomed to failure, and it is surprising that the Morlet transform in the discrete setting works as well as it does. While we have not given up entirely on the Morlet transform, if one wants to make quantitative use of the results of a spectral analysis, the energy and reconstruction theorems must be satisfied precisely.

How, then, are we to devise a multiresolution analysis that satisfies the fundamental theorems of spectral analysis? The answer lies in focusing our attention on the distribution of energy in the window function. Suppose we have one unit of energy, an “atom” so to speak, distributed with temporal power $\Phi_A^2(t)$ in the continuum, which we wish to use as the basis for a multiresolution analysis. Let the window at scale τ be defined by $\Phi_\tau(t) \propto \Phi_A(t/\tau)$ for positive integer τ , which scales the relative unit of time by integral amounts. One can supplement this definition for $\tau = 0$ by letting $\Phi_0(0) = 1$ and 0 otherwise. For an exponential decay with scale invariant parameter σ , one can write $\Phi_\tau(t) \propto \exp^{-\pi/2(t^2/\tau^2\sigma^2)}$, where the normalization to unit energy is done explicitly over the integers $t \in [-\tau, \tau]$. (Alternately, one could hold the window width fixed for all τ .) The layered window $\Phi_L(t)$ is then defined to be the sum in terms of energy, not amplitude, of the windows $\Phi_\tau(t)$ over some set of scales $\tau \in [\tau_1, \tau_L]$ with L members, normalized to unit energy, $\Phi_L^2(t) \equiv L^{-1} \sum_\tau \Phi_\tau^2(t)$. For a one-sided transform, $\Phi_L \rightarrow \sqrt{2} \Phi_L$ of course. This single window, with duration $T' = 2\tau_L + 1$, is to be used for all the frequency bins in the windowed Fourier transform.

Since the atomic window Φ_A is arbitrary, the distinction between summing the window energies Φ_τ^2 versus summing the window amplitudes Φ_τ is really just a matter of interpretation. Whichever way one defines the sum over scales, the layered window eventually is normalized explicitly to one unit of energy, or two units for a one-sided transform. Using the construction above in terms of energy, the amplitude of the layered window is the root-mean-square of the amplitudes of the scaled windows, while the construction summing the window amplitudes is not so simply normalized. With respect to its physical motivation, energy is the quantity that ultimately gets quantized, not amplitude, which is a property reflected in our favored definition of the layered window. The selection by the investigator of the set of scales $\{\tau\}$ to use in the construction specifies the domain of energy distributions to which the transform will be sensitive.

Let us return now to consideration of our four frequency test signal, with its mean restored. For comparison with the previous Gabor and Morlet transforms, let the decay parameter here be $\sigma = \sqrt{\pi}/6$, and let $\tau \in [12, 96]$. The results of such analysis are displayed in Figure ???. Compared with the Morlet transform, the spectral resolution is much improved in the upper portion of the frequency axis, as seen in panel (a). Since only one window is in operation, the marginal densities of panels (b) and (c) can be compared with the convolution of the window and signal energy densities in either the temporal or spectral representations. The absolute residual shown in panel (d) is on the order of the machine precision, with maybe a slight accumulation of truncation errors. The ratios $E_{\hat{y}}/E_y$ and $E_{\hat{g}}/E_{\hat{y}}$ equal unity to machine precision. The layered window transform, by virtue of using a single window for all frequencies, satisfies the energy and reconstruction theorems to the accuracy of one's computational device.

As a final comparison, let us put the Fourier power spectrum of the signal alongside the marginal spectral energy densities of the Gabor, Morlet, and layered window transforms, displayed in Figure ??. The Fourier spectrum in panel (a) is the typical mess one gets when analyzing non-stationary signals, but its net energy does equal the signal energy. The Gabor spectrum in panel (b) has lost its frequency resolution on account of the tight temporal bandwidth of the selected window parameters, yet retains the correct normalization. The Morlet spectrum in panel (c) resolves the low frequency portion of the spectrum but not the high frequency region and does not have the correct normalization. The

layered window spectrum in panel (d) is normalized to the signal energy and resolves both the low and high frequency parts of the spectrum while maintaining sufficient temporal resolution to be useful in identifying non-stationary features in the signal.

Figure 7. One-sided layered window transform of a real signal as described in the text. The power spectrum is shown in (a) with the signal’s instantaneous frequencies indicated by the dashed lines. The marginal spectral power is shown in (b) as \times , as is the marginal temporal power in (c), and each is compared with the convolution of the window and signal energies indicated by $+$. The absolute value of the reconstruction residual is displayed in (d).

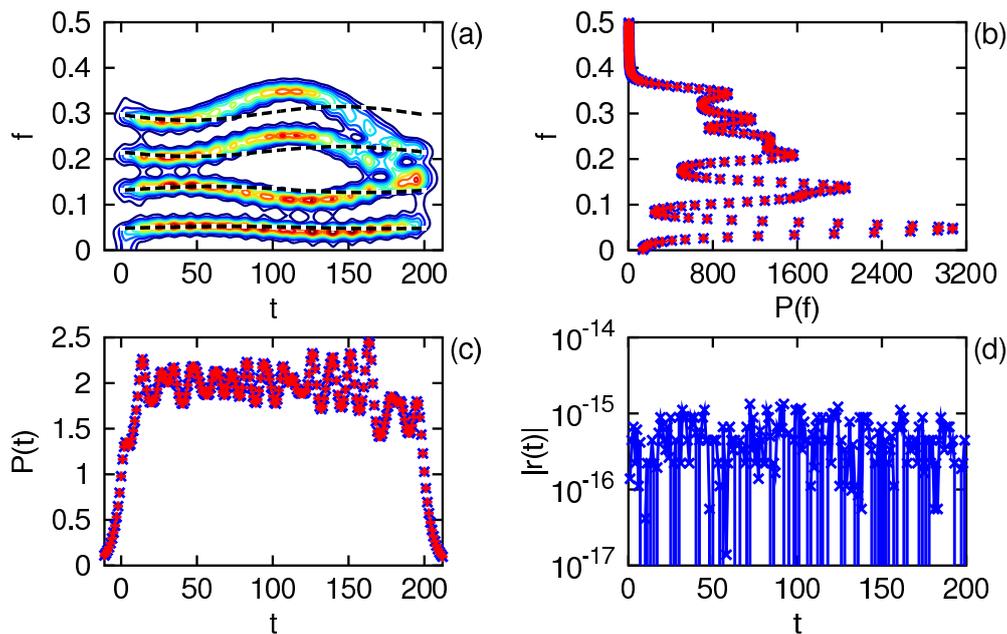
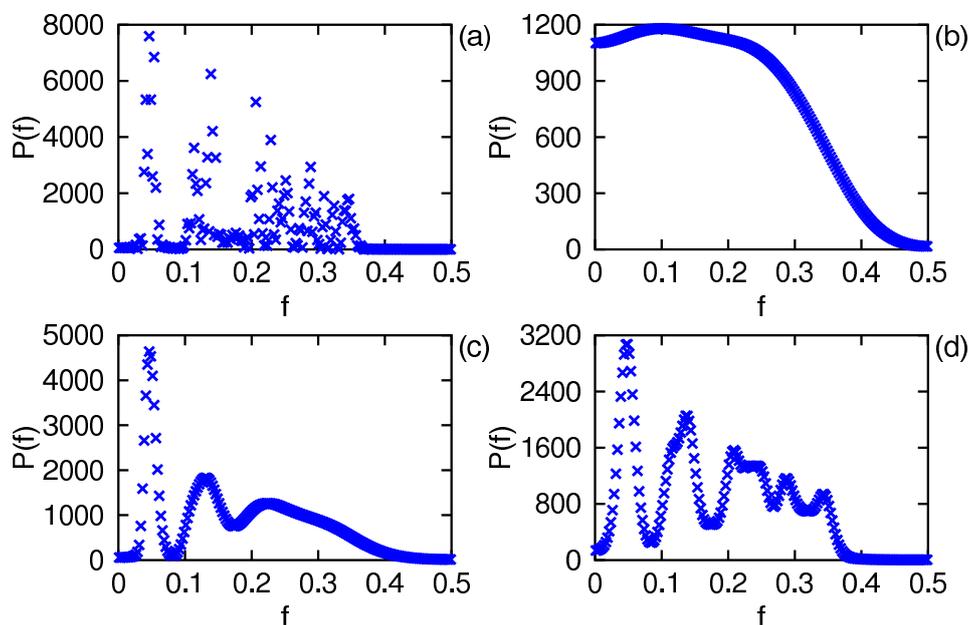


Figure 8. Comparison of the Fourier power spectrum in (a) with the marginal spectral energy densities of the Gabor transform in (b), the Morlet transform in (c), and the layered window transform in (d).



5. Irregular Sampling and Minimal Order

Let us now consider the case of an irregularly sampled signal $y(t)$, with regard to the analyses by [?] and [?]. Irregular sampling has long vexed practitioners of wavelet analysis, leading to a variety of suggestions for how to overcome its difficulties. Because of the nature of stellar observations, the astrophysical community often presents data that contains gaps or is otherwise on an irregular time axis. The method by [?] focuses on the normalization of the gapped wavelet basis functions, while the method by [?] focuses on the admissibility condition; taken together, those ideas lead to the edge adapted algorithm proposed by [?]. More familiar perhaps is the lifting scheme by [?], which operates in the context of the dyadic wavelet transform implemented in terms of finite impulse response digital filter banks.

For simplicity, we will assume that the samples each have a uniform duration Δ_t , which is not greater than any of the inter-measurement periods, so that the time metric D_t remains proportional to the identity matrix; otherwise, a suitably generalized D_t must be used. Let the observation times be given by the vector $\mathbf{t} \equiv t_d$ indexed by $d \in [1, D]$, where D is the total number of samples, in some unit u_t not necessarily equal to Δ_t such that the values t_d are integers; u_t is the resolution of the time measuring apparatus, thus Δ_t must also be an integer. Similarly, the measurements can be written as the vector $\mathbf{y} \equiv y_d$ so that the signal energy remains $E_y = \mathbf{y}^\dagger D_t \mathbf{y}$ in units of $u_y^2 u_t$. Missing values, *i.e.*, measurements at integer times not in \mathbf{t} , are effectively treated as zero. Let us look first at how the discretized Fourier transform is modified in this case.

To identify the Nyquist critical frequency, one must find the lowest positive Fourier basis function that is entirely real (technically up to a constant phase) over the given set of observation times [?]. That procedure is easily accomplished by shifting and scaling the time axis $\mathbf{t} \rightarrow \mathbf{t}'$ such that $f_c \equiv (2u_{t'})^{-1}$ in the new time units. If one defines t_g to be the greatest common divisor over the set of inter-measurement periods, the new time axis can be written $t'_d \equiv (t_d - t_1)/t_g + 1$ in units of $u_{t'} \equiv t_g u_t$ such that $t'_d \in [1, T]$ remains integer valued. At the critical frequency, the Fourier basis function is $\exp^{i\pi}(2f_c t'_d) = \pm 1$, and the Fourier spectrum has period 1 in units $u_f = u_{t'}^{-1}$. Having found the units in which $f_c = 1/2$, let us drop the prime distinguishing the scaled time axis in the following, and for convenience let us suppose $\Delta_t = 1$ in the scaled units; otherwise one must be extra careful with the scaling of the energy units.

The next task is to determine the minimal order for satisfaction of the energy and reconstruction theorems. Let $N_D \equiv \lceil D/2 \rceil$, and let $N_T \equiv \lceil T/2 \rceil$. For $N < N_D$ there are an insufficient number of DOF in the discrete Fourier transform to fully represent the information content of the (real) signal, thus the fundamental theorems are not satisfied in general. Likewise, for $N \geq N_T$ there are more than enough DOF to represent the signal; N_T gives the minimal order of the analogous regularly sampled signal where the missing entries are assigned the value zero. For $N_D \leq N < N_T$, the Fourier transform might work, depending upon the particulars of the irregular sampling. To satisfy the reconstruction theorem $\hat{\mathbf{y}} = \mathbf{y}$ for $\hat{\mathbf{y}} \equiv \mathbf{Q}\mathbf{y}$ in the one-sided case, one must have $\mathbf{Q} \equiv \text{Re } \Theta D_f \Theta^\dagger D_t = \mathbf{I}_D$ to the order of machine precision; the energy theorem also is satisfied when that condition is met. (For a complex signal, the entire matrix, not just its real part, must be utilized, and of course D_f must cover one full period of the frequency axis.) The basis functions are evaluated only at the observation times $\Theta(d, n) = \exp^{i2\pi}(f_n t_d)$,

because those are the only locations at which the reconstruction theorem must be satisfied. In Table ?? we compare the norm of the reconstruction residual $\mathbf{r} \equiv \hat{\mathbf{y}} - \mathbf{y}$ for some signal with $D = 10$ and $T = 28$ with the distance between the quality matrix \mathbf{Q} and the identity \mathbf{I}_D according to the Frobenius metric at various orders N with $P = 0$. To produce a smooth picture of the spectral content, one should take $N \gg N_T$.

Table 3. Comparison of the norm of the residual of the Fourier transform with the distance from the quality matrix \mathbf{Q} to the identity \mathbf{I}_D at various orders N and even parity for an irregularly sampled signal with $D = 10$ and $T = 28$ such that $N_D = 5$ and $N_T = 14$.

N	5	6	7	8	9
$\ \mathbf{r}\ _F$	3.5943e +00	5.1760e +00	3.0368e +00	5.5110e −15	3.0148e +00
$\ \mathbf{Q} - \mathbf{I}_D\ _F$	2.4495e +00	2.8284e +00	2.4495e +00	5.6505e −15	2.0000e +00
N	10	11	12	13	14
$\ \mathbf{r}\ _F$	6.6842e −01	5.4143e −15	2.6771e +00	6.2415e −15	7.4742e −15
$\ \mathbf{Q} - \mathbf{I}_D\ _F$	1.4142e +00	4.3648e −15	2.0000e +00	4.2983e −15	4.9648e −15

Turning now to the consideration of the Gabor transform (and by extension all windowed Fourier transforms with a fixed Φ), the construction of the quality matrix is a bit more complicated, owing to the convolutions along the time axis. It is commonly remarked that the number of DOF along the frequency axis is multiplied by the number of times the window duration $T' = 2\tau + 1$ fits within the signal duration T , but the situation is more subtle than that, as one must also account for the bandwidth of the window given by the decay σ . Beginning with the case of regular sampling $D = T$, using the phase convention of Equations (?? and ??), one can write the composition of the inverse and forward transforms as

$$\hat{y}(t) = \text{Re} \sum_{n=1}^{N'} \sum_{t'=-\tau}^{\tau} \Phi(t') \exp^{i2\pi(-f_n t')} \sum_{t''=-\tau}^{\tau} \Phi(-t'') \exp^{i2\pi(-f_n t'')} y(t+t'+t'') \Delta_t \Delta_t \Delta_f \tag{13a}$$

$$= \text{Re} \sum_{t'''=-2\tau}^{2\tau} \sum_{t'+t''=t'''} \Phi(t') \Phi(-t'') \sum_{n=1}^{N'} \exp^{i2\pi(-f_n t''')} y(t+t''') \Delta_f \Delta_t \Delta_t \tag{13b}$$

$$\equiv \sum_{t'''=-2\tau}^{2\tau} q(t''') y(t+t''') \Delta_t \tag{13c}$$

which yields one row in the quality matrix \mathbf{Q} indexed by $t''' \in [-2\tau, 2\tau]$ indicating the diagonal where $q(t''')$ appears such that \mathbf{Q} has the form of a Toeplitz matrix of rank T . For a symmetric window $\Phi(-t') = \Phi(t')$ written as a diagonal matrix Φ such that $\Psi \equiv \Phi\Theta$, the values $q(t''')$ can be extracted from $\Psi^* \mathbf{D}_f \Psi^\dagger$. For an irregularly sampled signal $T > D$, one retains only those rows and columns of \mathbf{Q} corresponding to actual data entries, $\mathbf{Q} \rightarrow \mathbf{Q}(\mathbf{t}, \mathbf{t})$. In Table ?? we compare the residual norm of the Gabor transform with $\tau = 12$ and two values of σ with the distance from \mathbf{Q} to \mathbf{I}_D at various N with $P = 0$ for an irregularly sampled signal of duration $T = 100$ with $D = 30$ entries. As N approaches N_D from below, the quality of reconstruction improves until the residual is on the order of machine precision.

Table 4. Comparison of the norm of the residual of the Gabor transform with width $\tau = 12$ and the indicated decay σ with the distance from the quality matrix Q to the identity I_D at various orders N and even parity for an irregularly sampled signal with $D = 30$ and $T = 100$ such that $N_D = 15$ and $N_T = 50$.

σ	N	6	7	8	9	10
$2\sqrt{\pi}$	$\ r\ _F$	5.4222e-04	2.8233e-05	5.1908e-07	1.1113e-08	7.3918e-11
	$\ Q - I_D\ _F$	5.2358e-04	2.0300e-05	4.4962e-07	7.0963e-09	5.1610e-11
	N	11	12	13	14	15
	$\ r\ _F$	2.7305e-13	1.2477e-14	6.2789e-15	5.8445e-15	7.2210e-15
	$\ Q - I_D\ _F$	1.6628e-13	4.3047e-16	4.1495e-16	2.8105e-16	2.2323e-16
	σ	N	6	7	8	9
$4\sqrt{\pi}$	$\ r\ _F$	4.5343e-01	2.6201e-01	7.5299e-02	3.4864e-02	6.9207e-03
	$\ Q - I_D\ _F$	4.3784e-01	1.8839e-01	6.5223e-02	2.2263e-02	4.8319e-03
	N	11	12	13	14	15
	$\ r\ _F$	1.0987e-03	7.8620e-05	1.0067e-14	1.2076e-14	1.2964e-14
	$\ Q - I_D\ _F$	6.6716e-04	6.9627e-05	6.0970e-16	7.7092e-16	1.0855e-15

For the Morlet basis functions, the construction of Ψ is more involved, as $\Psi(t', n) = \Phi(t', n)\Theta(t', n)$. Nonetheless, one can evaluate Q from $\Psi^*D_f\Psi^\dagger$ for any rank T . As a function of order N , one finds that the deviation $\|Q - I_T\|_F$ decreases until $N > T/2$, after which it bottoms out on the order of a few percent times T . In Table ?? we compare the norm of the residual for the Morlet basis using $\sigma = \sqrt{\pi}$ and $\tau = 6$ for some regularly sampled signal with $D = T = 30$ with the deviation of Q from I_D at various N with $P = 0$. If one wishes to investigate the feasibility of devising a Morlet basis with perfect reconstruction, the evaluation of Q is where to start.

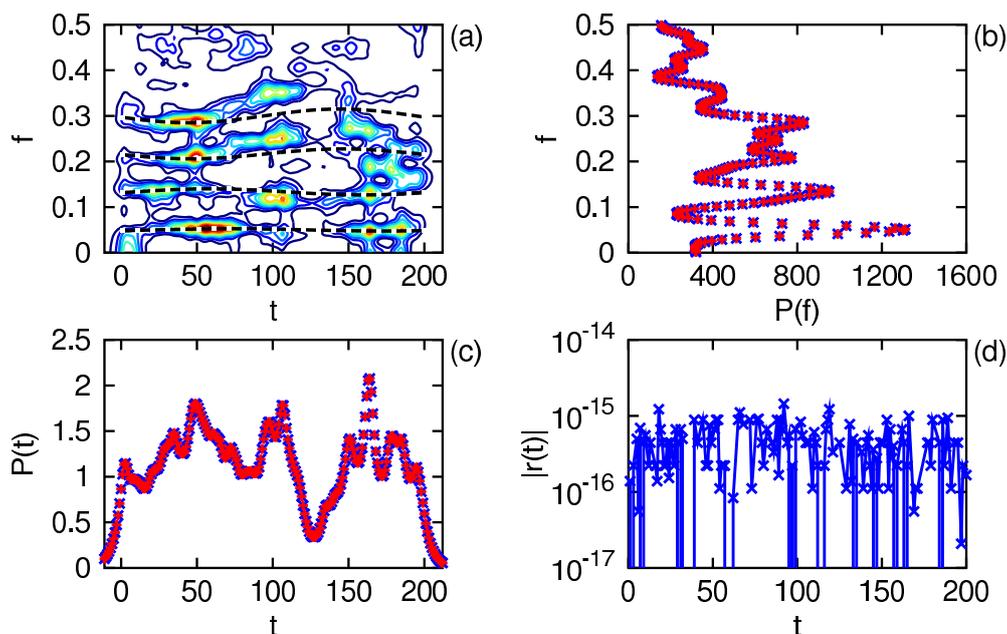
Table 5. Comparison of the norm of the residual of the Morlet transform with width $\tau = 6$ and decay $\sigma = \sqrt{\pi}$ with the distance from the quality matrix Q to the identity I_D at various orders N and even parity for a regularly sampled signal with $D = T = 30$ such that $N_D = N_T = 15$.

N	6	7	8	9	10
$\ r\ _F$	3.3471e+00	2.9397e+00	2.7428e+00	2.2582e+00	1.9475e+00
$\ Q - I_D\ _F$	2.9527e+00	2.4686e+00	2.0993e+00	1.8244e+00	1.5788e+00
N	11	12	13	14	15
$\ r\ _F$	1.7911e+00	1.5917e+00	1.3310e+00	9.3137e-01	5.3677e-01
$\ Q - I_D\ _F$	1.3451e+00	1.1139e+00	8.7648e-01	6.3175e-01	4.0199e-01

Let us close this section by looking at the analysis of the signal from the previous section but with the number of measurements reduced by a factor of a third ($N_D = 134$) for the same duration $T = 200$. The evaluation of the layered window Fourier transform proceeds as before, with the understanding that

values of $y(\hat{t} + t')$ at times not in \mathbf{t} are treated as zero; the time axis for the spectrum $\hat{Y}(n, \hat{t})$ includes every integer $\hat{t} \in [1 - \tau_L, T + \tau_L]$. Using the same parameters as before, $\sigma = \sqrt{\pi}/6$ and $\tau \in [12, 96]$ with $N = 200$ and $P = 0$, in Figure ?? we display the power spectrum, its marginal densities, and the reconstruction residual for our irregularly sampled signal. While the loss of information has affected the appearance of the spectrum relative to the regularly sampled case, it nonetheless remains a faithful representation of the available information. By virtue of using a single fixed window Φ_L , the layered window Fourier transform satisfies the fundamental theorems of spectral analysis even when there are gaps in the observation record.

Figure 9. One-sided layered window transform of an irregularly sampled signal as described in the text. The power spectrum is shown in (a) with the signal’s instantaneous frequencies indicated by the dashed lines. The marginal spectral power is shown in (b) as \times , as is the marginal temporal power in (c), and each is compared with the convolution of the window and signal energies indicated by $+$. The absolute value of the reconstruction residual is displayed in (d).



6. Window Comparison

Let us conclude the analysis with a comparison of the temporal and spectral bandwidths for several types of window function, as well as their estimates of the power spectral density carried by a test signal with a little more complexity than we had before. The test signal is regularly sampled with $T = 200$ and four component frequencies, but now only two components are present in the first half of the duration, while the other two are in the second half. An abrupt transition occurs at the midpoint of the duration. The amplitude of the frequency variation is now 10%, with a period of $T/2$, so that each component covers a broad range of instantaneous frequency. The order for all the transforms in this section will be the minimal Fourier order of the signal $N = 100$ with even parity $P = 0$.

The parameters for the four types of window considered are displayed in Table ??, comprising of two Gabor windows at either end of the range in scale for the layered window also considered, in addition to

a random window with an arbitrary duration. The first moments of time and frequency are indicated by t_1 and f_1 respectively, and the bandwidths (square root of the second moment about the first moment) by t_2 and f_2 . The two Gabor windows are observed to minimize the Fourier uncertainty relation $t_2 f_2 \geq 1/4\pi$, thus in that sense are optimal, while the layered window has a slightly larger bandwidth product. The random window has a bandwidth product that is considerably larger, yet all these windows produce a valid discrete transform pair that satisfies the fundamental theorems of spectral analysis as indicated by the residual of reconstruction $\|r\|_F$. The energy densities in the temporal and spectral representation used to evaluate the bandwidth products are shown in Figure ??, where one can see the trade-off in time/frequency resolution in action.

Table 6. Comparison of the temporal and spectral bandwidths for various windows, as well as their rms reconstruction residual for the test signal as described in the text.

Window	σ	τ	t_1	t_2	f_1	f_2	$4\pi t_2 f_2$	$\ r\ _F$
Gabor	$2\sqrt{\pi}$	12	−0.000	1.414	−0.000	0.056	1.000	0.000
Gabor	$16\sqrt{\pi}$	96	−0.000	11.314	0.000	0.007	1.000	0.000
Layered	$\sqrt{\pi}/6$	[12, 96]	0.000	6.990	0.000	0.014	1.244	0.000
Random	—	24	−2.836	13.718	0.000	0.130	22.442	0.000

The power spectral density estimates evaluated from the test signal of this section using the various windows are displayed in Figure ??, arranged from (a) to (d) according to Table ??. The power spectrum given by the layered window in (c) does indeed incorporate features of either Gabor transform shown in (a) and (b); in a sense, it has combined the Gabor transforms over its range of scale in a way that preserves the energy and reconstruction theorems. With some algebraic manipulation of the expressions defining the layered window Fourier transform one might be able to show that explicitly, but for now that statement is intuitive speculation. Interestingly, the spectrum given by the random window is not much different from the others, driving home the point that the shape of the window truly is arbitrary as long as it is normalized appropriately.

Figure 10. Comparison of the temporal and spectral energy densities of the various windows arranged according to Table ??.

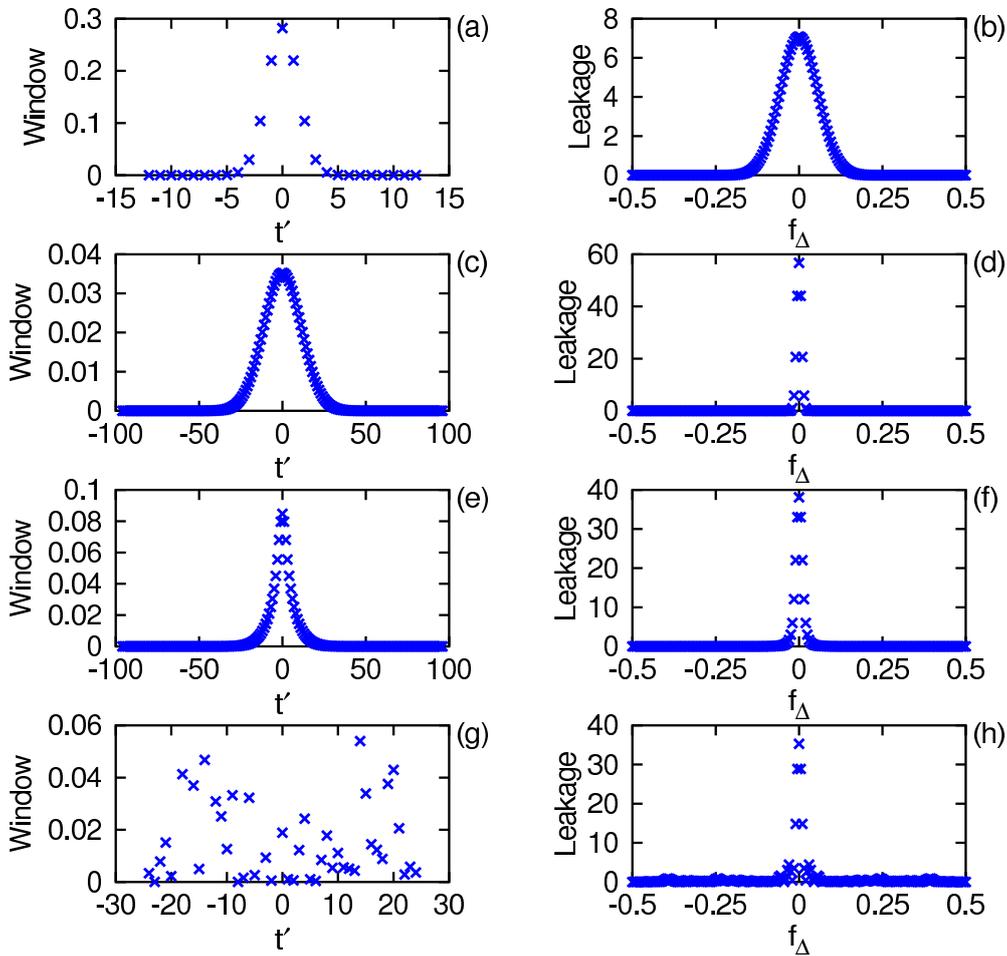
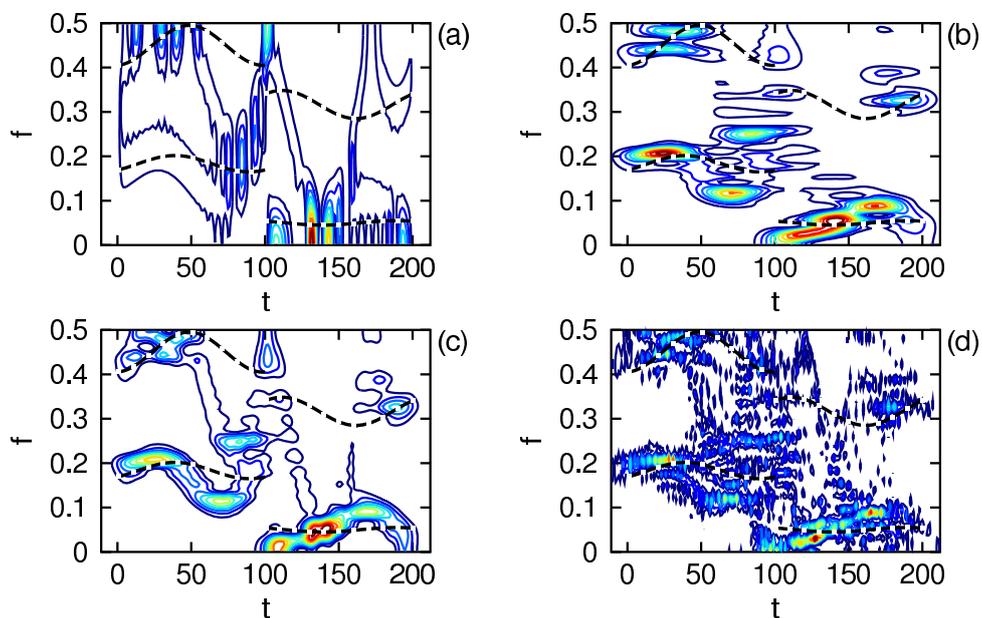


Figure 11. Comparison of the power spectral density of the test signal described in the text using the various windows arranged according to Table ?. The instantaneous frequencies used to generate the signal are indicated by the dashed lines.



7. Discussion

The primary result of this investigation is that the windowed Fourier transform has far more flexibility and utility than it is usually credited with. To achieve satisfaction of the energy and reconstruction theorems in the discrete setting, the only requirements on the window are that it be real, nonnegative, and normalized explicitly to unit energy, or two units for a one-sided transform. While we have looked only at symmetric windows here, one can easily verify that the expressions for the transform pair, either Equations (?? and ??) or (?? and ??), remain valid for a window that is not symmetric in t' ; with attention to the details of relocating the temporal bins, the coefficients can be assigned to the time corresponding to the peak of the window rather than its midpoint. Similarly, a window with an even length duration T' could be used if one is careful with the definition of the phase and the location of the bins. In fact, any nonnegative function can be used for Φ , or even a list of arbitrary numbers, as long as it is suitably normalized.

The flexibility in Φ allows one to define a multiresolution spectral analysis in terms of an atomic unit of energy $\Phi_A(t/\tau)$ evaluated over a range of scales τ . These multiple scalings of Φ_A are combined into a single layered window Φ_L that is applied to the entire frequency axis; the phase component of the basis is independent of the window function. Because only a single window is used, the fundamental theorems of spectral analysis are satisfied. With a Gaussian Φ_A , the form of Φ_L very closely resembles a Lorentzian function, which expresses uniformity over scale when expressed as an angle. The shape of Φ_L can be tailored to the needs of the investigation through inspection of its leakage function. Similarly to the wavelet transform, the trade-off between resolution in time or frequency is under the control of the investigator. Any structure seen in the spectral coefficients of the signal is understood to be conditioned on the selection of the window function; the answer one gets depends upon how the question is asked.

The difficulties faced by the continuous wavelet transform exemplified by the Morlet basis, with respect to the fundamental theorems of spectral analysis, can be summarized in the evaluation of its quality matrix Q . Perfect reconstruction in the discrete setting requires Q to equal the identity matrix to the precision of one's computational device. The various adjustments looked at above that improve the reconstruction residual must be bringing Q closer to that form; however, none of them achieve the required precision. For an integral transform pair to have quantitative significance, one must demonstrate that the distance from its Q to I is numerically zero. This investigation has proposed an alternate approach to multiresolution analysis, which does satisfy the energy and reconstruction theorems while improving upon the resolution properties of the Gabor transform.

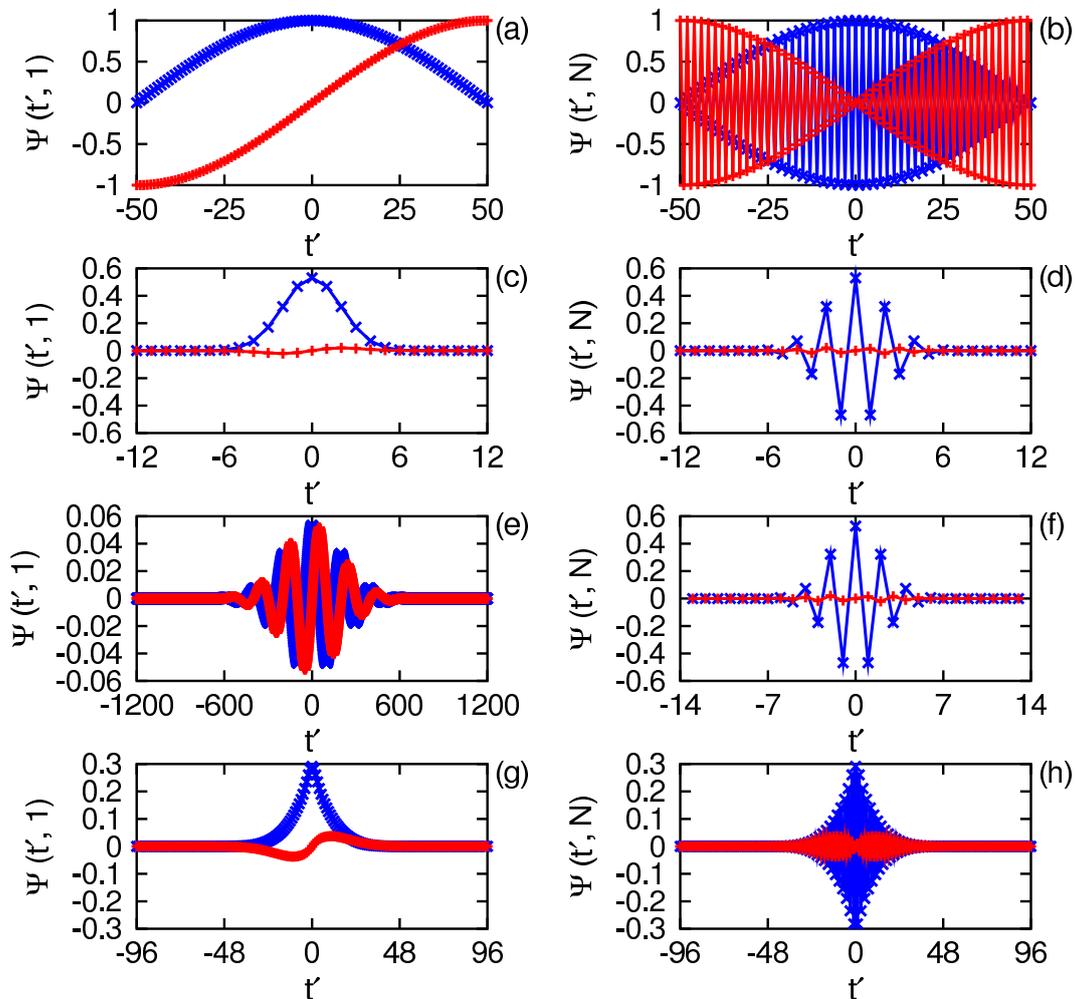
The effects of aliasing and of irregular sampling are easily understood with regard to the periodicity induced in the Fourier spectrum. The Nyquist interval is given simply by the span between frequencies that are indistinguishable over the stated measurement times; assigning a value of unity to that range such that $f_c = 1/2$ is just a matter of choosing the appropriate units for time. The implementation of the windowed Fourier transform is unaffected beyond zero padding the signal at locations of missing measurements, whose justification is that no other procedure leaves the energy, hence information content, unchanged. One should observe that the ends of a regularly sampled signal are treated the same way. On that note, we can remark that there is no cone of influence for the windowed Fourier

transform, as every spectral coefficient is important to the satisfaction of the fundamental theorems, even those outside the temporal domain of the data.

The minimal order N_{\min} required for a faithful representation of the signal depends upon the temporal bandwidth of the window and the duration of the data. Pinning that number down for the windowed Fourier transform is a bit tricky, and if needed is best evaluated explicitly through inspection of Q as a function of order N . What one can say in general is that $1 \leq N_{\min} \leq N_T$, where the bounds are given by the minimal orders for the fully temporal and fully spectral representations of the regularly sampled signal, respectively. Irregular sampling complicates matters by introducing an order $N_D < N_T$ such that $N_D \leq N_{\min} \leq N_T$ for the full-length Fourier transform and $N_{\min} \leq N_D$ for the windowed Fourier transform. For most practical cases of data analysis, one is interested in producing a smooth plot of the spectral content, thus one would use an order $N \gg N_{\min}$. To verify the marginal density of the windowed power spectrum in comparison with the convolution of the window and signal energy densities in the spectral representation, one requires an order N sufficient to resolve both the signal and the window temporal durations.

Thus far we have heard nary a peep from the actual basis functions used in this analysis, so let us close the discussion by looking at some of the stars of the show. In Figure ?? we display the real and imaginary parts of the basis functions, normalized to unit energy, for each of the transforms considered above at order $N = 50$ with even parity $P = 0$ at the lowest and highest entries on the frequency axis, *i.e.*, $f_1 = 1/200$ and $f_N = 99/200$. The Fourier transform has no intrinsic window, so it is shown for a duration $T' = 101$ that spans $T < T'$. The window parameters chosen are those from the previous sections: $\sigma = 2\sqrt{\pi}$ and $\tau = 12$ for the Gabor transform, $\sigma = \sqrt{\pi}$ and $\tau = 6$ for the Morlet transform, and $\sigma = \sqrt{\pi}/6$ and $\tau \in [12, 96]$ for the layered window transform. One can observe that the transforms whose Q equals the identity share the property that these basis functions are phase analogues over the discrete sample times, whereas the Morlet basis functions are scale analogues. By phase analogue we mean if $\Psi(t', 1) = a_1 + ib_1$, then $\Psi(t', N) = a_1(-1)^{t'} + ib_1(-1)^{t'+1}$, a feature shared by all positive frequency pairs whose midpoint is $f = 1/4$; the Morlet basis differs from the others in this regard. Whether that property is required for satisfaction of the fundamental theorems is unclear, but the results of this analysis suggest that it is.

Figure 12. Comparison of the lowest $\Psi(t', 1)$ and highest $\Psi(t', N)$ frequency basis functions at order $N = 50$ with $P = 0$ using the parameters found in the text, with the real part indicated by \times and the imaginary part by $+$. The Fourier basis is in (a) and (b), the Gabor basis is in (c) and (d), the Morlet basis is in (e) and (f), and the layered window basis is in (g) and (h).



8. Conclusions

In this article we have compared the Fourier, Gabor, and Morlet transforms in the fully discrete setting applicable to the analysis of sampled data. While the Fourier and Gabor transforms satisfy the fundamental theorems of energy conservation and perfect reconstruction, the Morlet transform does not. The magnitude of the residual can be related to the distance from the quality matrix to the identity in all cases, such that the minimal order for reconstruction can be determined for the Fourier and Gabor transforms. Various methods of improving the response of the Morlet transform are considered; however, none of them achieve the desired precision on par with the truncation error of one’s computational device.

An alternate approach to multiresolution analysis is proposed, which constructs a single layered window from multiple scalings of some atomic unit of energy. This layered window transform satisfies the fundamental theorems of spectral analysis while providing sufficient temporal resolution to identify non-stationary features in the signal. The trade-off between time and frequency resolution is under

the control of the investigator through the selection of the atomic window and the scales over which it is evaluated. The power spectrum of the layered window transform is similar to that of the Morlet transform but provides much better frequency resolution.

The premise behind the wavelet transform is that the low frequency elements of a signal should have a much longer duration than the high frequency elements. There are, however, many examples of real world signals for which the converse is true. Consider, for example, the digital recording of a kick drum and cymbal rhythm such that the low frequency bursts have a relatively short duration compared with the ringing at high frequency. For that type of signal, the wavelet transform is going to provide a poor choice of basis even if it satisfied the fundamental theorems. The flexibility in the choice of window function used in the windowed Fourier transform allows one to tailor its response to the needs of the analysis far better. The result one gets depends upon how one defines the resolution of the window, any of which provide a valid spectral representation of the signal.

To make quantitative use of the spectral analysis of some discretely sampled signal, one must demonstrate that the fundamental theorems are satisfied. The windowed Fourier transform can be shown to satisfy those requirements for any real valued window function that is suitably normalized. The potential for gaps in the data, or some other form of irregular sampling, is found to pose no problem once the correct Nyquist interval is identified. Consequently, the windowed Fourier transform can be applied to data from a wide variety of sources, such as astronomical observations, which are limited physically to an irregular set of observation times, as well as the common case of regular sampling. The implementation of a multiresolution spectral analysis in the discrete setting appears to be possible only when the multiple scalings of the energy distribution are applied evenly across the frequency axis by combining them into a single layered window.

References

- . Torrence, C.; Compo, G.P. A practical guide to wavelet analysis. *Bull. Am. Meteor. Soc.* **1998**, *79*, 61–78.
- . Farge, M. Wavelet transforms and their applications to turbulence. *Ann. Rev. Fluid Mech.* **1992**, *24*, 395–458.
- . Grossmann, A.; Morlet, J. Decomposition of hardy functions into square integrable wavelets of constant shape. *SIAM J. Math. Anal.* **1984**, *15*, 723–736.
- . Meyer, Y. Ondelettes et fonctions splines. Available online: <https://eudml.org/doc/111928> (accessed on 19 June 2013).
- . Mallat, S. A theory for multiresolution signal decomposition: The wavelet representation. *Pattern Anal. Mach. Intell. IEEE Trans.* **1989**, *11*, 674–693.
- . Daubechies, I. Orthonormal bases of compactly supported wavelets. *Commun. Pure Appl. Math.* **1988**, *41*, 909–996.
- . Vetterli, M.; Herley, C. Wavelets and filter banks: Theory and design. *Signal Process. IEEE Trans.* **1992**, *40*, 2207–2232.
- . Kronland-Martinet, R.; Morlet, J.; Grossmann, A. Analysis of sound patterns through wavelet transforms. *Int. J. Pattern Recognit. Artif. Intell.* **1987**, *1*, 273–302.

- . Meyers, S.D.; Kelly, B.G.; O'Brien, J.J. An introduction to wavelet analysis in oceanography and meteorology: With application to the dispersion of yanai waves. *Mon. Weather Rev.* **1993**, *121*, 2858–2866.
- . Baliunas, S.; Frick, P.; Sokoloff, D.; Soon, W. Time scales and trends in the central England temperature data (1659–1990): A wavelet analysis. *Geophys. Res. Lett.* **1997**, *24*, 1351–1354.
- . Fligge, M.; Solanki, S.K.; Beer, J. Determination of solar cycle length variations using the continuous wavelet transform. *Astron. Astrophys.* **1999**, *346*, 313–321.
- . Christopoulou, E.B.; Skodras, A.N.; Georgakilas, A.A. Time Series Analysis of Sunspot Oscillations Using the Wavelet Transform. In Proceedings of the 14th International Conference on Digital Signal Processing, Santorini, Greece, 1–3 July, 2002; Volume 2; pp. 893–896.
- . Le, G.M.; Wang, J.L. Wavelet analysis of several important periodic properties in the relative sunspot numbers. *Chin. J. Astron. Astrophys.* **2003**, *3*, 391–394.
- . Grinsted, A.; Moore, J.C.; Jevrejeva, S. Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Process. Geophys.* **2004**, *11*, 561–566.
- . Piscaronoft, P.; Kalvová, J.; Brázdil, R. Cycles and trends in the Czech temperature series using wavelet transforms. *Int. J. Climatol.* **2004**, *24*, 1661–1670.
- . Liu, Y.; Liang, X.S.; Weisberg, R.H. Rectification of the bias in the wavelet power spectrum. *J. Atmos. Ocean. Technol.* **2007**, *24*, 2093–2102.
- . Jevrejeva, S.; Moore, J.C.; Grinsted, A. Influence of the Arctic Oscillation and El Niño-Southern Oscillation (ENSO) on ice conditions in the Baltic Sea: The wavelet approach. *J. Geophys. Res. Atmos.* **2003**, *108*, doi:10.1029/2003JD003417.
- . Echer, M.P.S.; Echer, E.; Nordemann, D.J.R.; Rigozo, N.R. Multi-resolution analysis of global surface air temperature and solar activity relationship. *J. Atmos. Solar-Terr. Phys.* **2009**, *71*, 41–44.
- . Greene, N. Inverse wavelet reconstruction for resolving the Gibbs phenomenon. *Int. J. Circuits Syst. Signal Process.* **2008**, *2*, 73–77.
- . Alphawave Toolbox for Spectral Analysis. Available online: <http://www.alphawaveresearch.com> (accessed on 20 May 2013).
- . Press, W.; Teukolsky, S.; Vetterling, W.; Flannery, B. *Numerical Recipes in C*, 2nd ed.; Cambridge University Press: Cambridge, UK, 1992.
- . Sadowsky, J. Investigation of signal characteristics using the continuous wavelet transform. *Johns Hopkins APL Tech. Dig.* **1996**, *17*, 258–269.
- . Johnson, R.W. *Wavelets: Classification, Theory and Applications*; Nova Science Publishers: Hauppauge, NY, USA, 2012; Chapter 6, pp. 125–155.
- . Johnson, R.W. Symmetrization and enhancement of the continuous Morlet transform for spectral density estimation. *Int. J. Wavelets Multiresolut. Inf. Process.* **2012**, *10*, doi:10.1142/S0219691311004511.
- . Lomb, N.R. Least-squares frequency analysis of unequally spaced data. *Astrophys. Space Sci.* **1976**, *39*, 447–462.
- . Scargle, J.D. Studies in astronomical time series analysis. II. Statistical aspects of spectral analysis of unevenly spaced data. *Astrophys. J.* **1982**, *263*, 835–853.

- . Foster, G. Wavelets for period analysis of unevenly samples time series. *Astron. J.* **1996**, *112*, 1709–1729.
- . Frick, P.; Baliunas, S.L.; Galyagin, D.; Sokoloff, D.; Soon, W. Wavelet analysis of stellar chromospheric activity variations. *Astrophys. J.* **1997**, *483*, 426–434.
- . Johnson, R.W. Edge adapted wavelets, solar magnetic activity, and climate change. *Astrophys. Space Sci.* **2010**, *326*, 181–189.
- . Sweldens, W. The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.* **1998**, *29*, 511–546.
- . Johnson, R.W. MaxEnt power spectrum estimation using the Fourier transform for irregularly sampled data applied to a record of stellar luminosity. *Astrophys. Space Sci.* **2012**, *338*, 35–48.

© 2013 by the author; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).