

INFORMATION TO USERS

While the most advanced technology has been used to photograph and reproduce this manuscript, the quality of the reproduction is heavily dependent upon the quality of the material submitted. For example:

- Manuscript pages may have indistinct print. In such cases, the best available copy has been filmed.
- Manuscripts may not always be complete. In such cases, a note will indicate that it is not possible to obtain missing pages.
- Copyrighted material may have been removed from the manuscript. In such cases, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, and charts) are photographed by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each oversize page is also filmed as one exposure and is available, for an additional charge, as a standard 35mm slide or as a 17"x 23" black and white photographic print.

Most photographs reproduce acceptably on positive microfilm or microfiche but lack the clarity on xerographic copies made from the microfilm. For an additional charge, 35mm slides of 6"x 9" black and white photographic prints are available for any photographs or illustrations that cannot be reproduced satisfactorily by xerography.



Order Number 8718732

A high resolution data-adaptive time-frequency representation

Jones, Douglas Llewellyn, Ph.D.

Rice University, 1987

U·M·I
300 N. Zeeb Rd.
Ann Arbor, MI 48106



PLEASE NOTE:

In all cases this material has been filmed in the best possible way from the available copy. Problems encountered with this document have been identified here with a check mark .

1. Glossy photographs or pages _____
2. Colored illustrations, paper or print _____
3. Photographs with dark background _____
4. Illustrations are poor copy _____
5. Pages with black marks, not original copy _____
6. Print shows through as there is text on both sides of page _____
7. Indistinct, broken or small print on several pages
8. Print exceeds margin requirements _____
9. Tightly bound copy with print lost in spine _____
10. Computer printout pages with indistinct print _____
11. Page(s) _____ lacking when material received, and not available from school or author.
12. Page(s) _____ seem to be missing in numbering only as text follows.
13. Two pages numbered _____. Text follows.
14. Curling and wrinkled pages
15. Dissertation contains pages with print at a slant, filmed as received
16. Other _____

University
Microfilms
International



RICE UNIVERSITY

A HIGH RESOLUTION DATA-ADAPTIVE TIME-FREQUENCY REPRESENTATION


by

DOUGLAS LLEWELLYN JONES

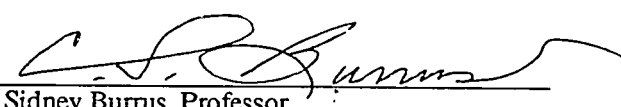
A THESIS SUBMITTED
IN PARTIAL FULFILLMENT OF THE
REQUIREMENTS FOR THE DEGREE

DOCTOR OF PHILOSOPHY

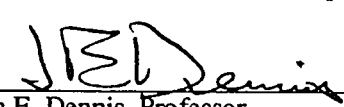
APPROVED, THESIS COMMITTEE:



Thomas W. Parks, Professor
Department of Electrical & Computer Engineering
Director



C. Sidney Burrus, Professor
Department of Electrical & Computer Engineering



John E. Dennis, Professor
Department of Mathematical Sciences

Houston, Texas

April 1987

ABSTRACT

A HIGH RESOLUTION DATA-ADAPTIVE TIME-FREQUENCY REPRESENTATION

by

Douglas Llewellyn Jones

The short-time Fourier transform and the Wigner distribution are the time-frequency representations that have received the most attention. The Wigner distribution has a number of desirable properties, but it introduces nonlinearities called cross-terms that make it difficult to interpret when applied to real multi-component signals. The short-time Fourier transform has achieved widespread use in applications, but it often has poor resolution of signal components and can bias the estimate of signal parameters. A need exists for a time-frequency representation without the shortcomings of the current techniques.

This dissertation develops a data-adaptive time-frequency representation that overcomes the often poor resolution of the traditional short-time Fourier transform, while avoiding the nonlinearities that make the Wigner distribution and other bilinear representations difficult to interpret and use. The new method uses an adaptive Gaussian basis, with the basis parameters varying at different time-frequency locations to maximize the local signal concentration in time-frequency. Two methods for selecting the Gaussian parameters are presented: a method that maximizes a measure of local signal concentration, and a parameter estimation approach. The new representation provides much better performance than any of the currently known techniques in the analysis of multi-modal dispersive waveforms.

ACKNOWLEDGMENTS

I would like to gratefully acknowledge the support that a number of persons have given me throughout this project. I would first like to thank my advisor, Dr. T.W. Parks, for his guidance and support. Special thanks are due to Dr. C.S. Burrus for watching out for me and for his counsel. I thank H.V. Sorensen, D. Linebarger, M.T. Heideman, S. Cabrera, S. McCaslin, J.W. Clark, J.E. Dennis, D.H. Johnson, C. Ganesh, M. Dahleh, D. Williams, V. Benson, N. Quioco, H. Daugherty, D. Schroeder, and all the other graduate students, faculty, and staff who have contributed both professional and personal support during the course of this work. Special thanks are due to G.L. Duckworth and A.B. Baggeroer for supplying me with data and with comments on the application of time-frequency analysis procedures to Arctic underwater acoustic data. I also want to thank my parents for their constant encouragement and support, and my wife Kitty for her patience, encouragement, and support, and also for proofreading and helping with the figures.

I also gratefully acknowledge the financial support of a National Science Foundation Graduate Fellowship, NSF grant ECS 83-14006, and Schlumberger Well Services. This work could not have been done without their support.

Table of Contents

1	Introduction	1
2	Current Time-Frequency Representations	3
	2.1 Time-Frequency Uncertainty, Concentration, and Resolution	3
	2.2 The Short-Time Fourier Transform	4
	2.3 Window Effects in the Short-Time Fourier Transform	5
	2.4 Short-Time Spectral Estimates	7
	2.5 The Wigner Distribution	8
	2.6 Other Time-Frequency Representations	10
3	A Data-Adaptive Time-Frequency Representation	12
4	Optimal Window Parameter Selection by Maximizing Concentration	16
	4.1 Introduction	16
	4.2 A Measure of Time-Frequency Signal Concentration	16
	4.3 A Measure of Local Time-Frequency Signal Concentration	19
	4.4 Implementation of the New Method	22
	4.5 Extensions to Improve Computation Speed	24
5	Optimal Window Selection by Gaussian Parameter Estimation	28
	5.1 Introduction	28
	5.2 A Gaussian Parameter Estimate	29
	5.3 A Modified Gaussian Parameter Estimate With Frequency Selectivity	30
	5.4 Implementation of the Parameter Estimation Method	30
6	Applications and Comparisons	33
	6.1 Introduction	33
	6.2 Applications of Time-Frequency Analysis	33
	6.3 Application to Synthetic Whistler Analysis	36
	6.4 Analysis of Two Crossing Gaussian Chirps	41
	6.5 Analysis of a Synthetic Arctic Acoustic Signal	51
	6.6 Analysis of a Real Arctic Acoustic Signal	65
7	Conclusion	71
	7.1 Conclusions	71
	7.2 Suggestions for Further Research	71
	References	74

CHAPTER 1

Introduction

Time-frequency representations describe signals in terms of their frequency content at a given time. These representations have proven to be useful for analyzing signals with both time and frequency variations, such as speech, music, and certain propagating waves. Of particular interest are signals propagating in dispersive channels such as the borehole complex in acoustic well-logging, the Arctic underwater acoustic channel, and the mantle of the earth.

The short-time Fourier transform has been widely used in time-frequency analysis for a number of years. More recently, bilinear time-frequency representations (most notably, the Wigner distribution) have been developed or revived as possible high-resolution alternatives to the short-time Fourier transform. The Wigner distribution has a number of properties, including high signal concentration in time-frequency, that are very desirable in a time-frequency representation. However, the Wigner distribution is highly nonlinear; it creates cross-terms between multiple signal components in time-frequency. These cross-terms can obscure other components or even closely-spaced auto-components, and the cross-terms can also be mistaken for true signal components. Only the short-time Fourier transform has achieved widespread use in applications, even though it is considered to have less resolution than the bilinear representations, because the bilinear representations are so difficult to interpret when applied to realistic multi-component signals. A need exists for a high resolution time-frequency representation that avoids the cross-terms associated with the Wigner distribution and other bilinear representations.

This thesis develops a linear, high resolution, data-adaptive time-frequency representation based on Gaussian "windows." This new representation is somewhat similar to the short-time Fourier transform, except that the "window" function varies with time and frequency to maximize the time-frequency concentration of the locally dominant component. The second chapter of this document describes the current time-frequency representations and their behavior. The third chapter presents the central concept of the new data-adaptive time-frequency representation and the reasons why such a representation can have superior time-frequency resolution and signal concentration. A local measure of time-frequency signal concentration is defined in the fourth chapter, and a procedure for determining the optimal Gaussian window parameters that maximize this concentration measure is developed. This procedure requires a large amount of computation, so an alternate method based on a Gaussian parameter estimate is presented in the fifth chapter. This technique requires much less computation than the concentration-based method but does not perform as well. The sixth chapter presents a number of examples that demonstrate the performance of the new time-frequency representations. The final chapter contains conclusions and some suggestions for further research.

CHAPTER 2

Current Time-Frequency Representations

2.1. Time-Frequency Uncertainty, Concentration, and Resolution

A time-frequency representation describes a signal in terms of its joint time and frequency content. A time-frequency representation is in some sense inherently ill-posed mathematically, since a time-frequency representation results in an artificial increase in the dimension of a signal by mapping a signal that is a function of one variable (time) to a function of two variables (time and frequency). This artificial increase in dimension implies that not all functions in the time-frequency plane have a corresponding time function. The manifold of valid time-frequency representations is difficult to characterize, but several results are known. Probably the most important result is the time-frequency uncertainty principle, which gives an upper bound for the concentration of any signal in the time-frequency plane. This result is analogous to Heisenberg's uncertainty principle in quantum mechanics; the time and frequency content of a signal are not independent, and the joint time-frequency content of a signal can be determined only up to the uncertainty limit.

Maximum concentration of signal components in time-frequency is often a goal of time-frequency analysis, since high concentration usually results in good resolution. Two signal components are considered resolved in time-frequency if their peaks are both visible as local maxima. In time-varying filtering applications using an analysis-synthesis technique, maximum concentration minimizes the area in time-frequency occupied by a signal component, thus allowing a tighter filter with greater rejection of noise and unwanted signal components. Several definitions of time-frequency concentration have been developed by various

researchers [15, 19, 32]. The development of new time-frequency representations is largely a quest for useful representations that approach the maximum concentration allowed by the time-frequency uncertainty principle; this is the primary goal of this thesis as well. In this chapter we describe the time-frequency representations that are currently of interest, and the qualities and drawbacks of each of these methods are discussed.

2.2. The Short-Time Fourier Transform

The short-time Fourier transform, defined as

$$S(t, \omega) = \int_{-\infty}^{\infty} x(\tau)w(\tau-t)e^{-j\omega\tau}d\tau, \quad (2.1)$$

in which $x(t)$ is the signal and $w(t)$ is the window function, has been widely used in time-frequency analysis for a number of years. This time-frequency representation has traditionally been viewed in two different ways. The short-time Fourier transform can be thought of as a bank of band-pass filters with impulse responses $w(-t)e^{j\omega t}$; the demodulated output of one of the bandpass filters represents the signal content at that particular frequency and time. Alternatively, $w(t)$ can be considered a window that selects a particular portion of the signal centered around the given time location, and the Fourier transform of the windowed signal yields the frequency content of the signal at the given time.

A third viewpoint yields insight into the short-time Fourier transform that is fundamental to the new methods developed in this thesis. The function $w(\tau-t)e^{-j\omega\tau}$ is concentrated in time-frequency around the location (t, ω) . The short-time Fourier transform is a projection of the signal onto a basis formed by a set of these functions; the projection, or inner product, with a particular element is a measure of the time-frequency content of the signal at that location. Gabor first applied this point of view to time-frequency analysis in 1946 [15].

Ideally, the projection function, or modulated window, should be an impulse in time-frequency. However, Gabor proved the time-frequency uncertainty principle, which precludes the existence of impulses in time-frequency. Gabor found that Gaussian signals $e^{-j\omega_0(t-t_0)+c(t-t_0)^2}$ achieve minimum time-frequency uncertainty, which implies that they are the closest approximation to an impulse in time-frequency; hence, time-shifted and frequency-modulated Gaussian functions appear to be the best basis in a projection-based time-frequency representation such as the short-time Fourier transform. The short-time Fourier transform with a Gaussian window is sometimes called a Gabor-Helstrom transform, since Helstrom came up with this representation by generalizing Gabor's results [16]. Other researchers have applied different measures of time-frequency concentration and have found that Gaussian functions maximize these measures as well [19, 32].

2.3. Window Effects in the Short-Time Fourier Transform

Many researchers have found that the choice of the window dramatically affects the appearance of and the signal concentration in the short-time Fourier transform. The optimality of the matched filter ($w(t) = x_c^*(t)$) is well known for detection purposes, where $x_c(t)$ is the signal component of interest. It is also well known in radar applications that tradeoffs can be made in time-frequency resolution by varying the length and time-frequency orientation of the window to enhance resolution in one direction in time-frequency at the expense of resolution in the orthogonal direction [28].

Barber and Stegun have determined the optimal bandwidth of the window for signals with a time-varying frequency to be roughly equal to the square root of the time derivative of the instantaneous frequency of the signal [1]. For some signals, this leads to the use of analysis filters in a generalization of the short-time Fourier transform with different bandwidths at different frequencies [31]. A number of authors have suggested the use of

frequency-varying filters in time-frequency representations [13, 25, 30]. These analyses, however, neglect the orientation of the filter in time-frequency. Dudgeon has demonstrated the value of matching the chirp rate of a window to that of the signal [12]; the chirp rate of a window is another degree of freedom besides the bandwidth that can be exploited to improve the performance of the short-time Fourier transform. More recently, it has been shown that the window maximizing the concentration of localized signal components in time-frequency is the matched window ($w(t) = x_c^*(-t)$), and that the effective resolution of the matched signal component in the short-time Fourier transform is at least equal to that in the Wigner distribution [21].

The results of the past research into the effects of windows in the short-time Fourier transform can be summarized as follows:

- (1) The window function has a major effect on the concentration and resolution in a short-time Fourier transform.
- (2) The best window depends on the signal, and may differ for different signal components in the same signal.
- (3) Use of a window that varies with frequency results in better performance on certain types of signals.
- (4) A window that is matched in some way to the signal component yields high performance according to several measures.
- (5) Gaussian windows are optimal according to several measures of time-frequency uncertainty and concentration.

These observations lead to the idea for a new time-frequency representation, to be described in Chapter 3 of this dissertation.

2.4. Short-Time Spectral Estimates

The short-time Fourier transform at a fixed time location is a spectral estimate of a segment of the signal. Several authors have considered applying high resolution spectral estimates instead of the DFT with a window to obtain better time-frequency resolution [9, 24]. Although these spectral estimates have somewhat better frequency resolution than the DFT, in many cases they yield a short-time spectral representation that is much inferior to the short-time Fourier transform with a properly chosen window. These approaches are mostly model-based methods that assume stationarity of the signal, but the purpose of time-frequency analysis is to process signals with frequency content that varies with time. The fundamental signal models on which these approaches are based are thus inappropriate.

The high resolution spectral estimates fail to process chirp signals effectively. A chirp signal occupies a large portion of the frequency spectrum over the analysis interval. A window with a chirp rate matched to that of the signal demodulates the chirp in the signal, so the windowed signal occupies a narrow frequency band, resulting in a spectrum occupying a much smaller frequency band than the signal in the analysis interval. The high resolution spectral estimates do not contain this implicit chirp demodulation, so they produce a spectrum occupying the full bandwidth of the signal in the analysis interval, resulting in spectral estimates with much less resolution than the DFT with the chirp demodulation in the window.

The discussion in the previous section indicates that the best length of the analysis interval depends on the signal and that this length may vary with different components of a

multi-component signal. The analysis interval must be as carefully chosen for a high resolution spectral analysis method as for the DFT and may be acceptable only for a single signal component. The high resolution spectral estimates thus suffer from the same drawbacks as the traditional short-time Fourier transform and are furthermore unsuitable for processing signals with chirp components. For certain types of signals they can provide a modest improvement over the traditional short-time Fourier transform, but they are not a sound general-purpose alternative to the short-time Fourier transform.

2.5. The Wigner Distribution

The Wigner distribution of a signal,

$$WD_x(t, \omega) = \int_{-\infty}^{\infty} x\left(t + \frac{\tau}{2}\right) x^*\left(t - \frac{\tau}{2}\right) e^{-j\omega\tau} d\tau, \quad (2.2)$$

is of special theoretical importance in time-frequency analysis [4-7, 19, 34]. The Wigner distribution preserves the time and the frequency marginals of a signal, or

$$|x(t)|^2 = \frac{1}{2\pi} \int_{-\infty}^{\infty} WD_x(t, \omega) d\omega, \quad (2.3)$$

$$|X(\omega)|^2 = \int_{-\infty}^{\infty} WD_x(t, \omega) dt. \quad (2.4)$$

The Wigner distribution also preserves the convex hull of the support in time-frequency. In other words, if the signal is zero outside a connected band in time or in frequency, the Wigner distribution is also zero for times or frequencies outside this band.

The above properties are quite desirable in a time-frequency representation, because they intuitively suggest that the Wigner distribution places signal energy in the proper place in the time-frequency plane, that the Wigner distribution achieves the time-frequency uncer-

tainty bound in terms of signal concentration, and that the Wigner distribution may be interpreted as a time-frequency signal energy distribution. However, the Wigner distribution of a signal in general will be negative in some regions of the time-frequency plane, so an energy distribution interpretation is not strictly correct.

A more serious problem with the Wigner distribution is the presence of cross-terms introduced by the nonlinearity of the Wigner distribution. The Wigner distribution of the sum of two signals is the sum of the Wigner distributions of the two signals plus an interference term between the signals;

$$WD_{x+y} = WD_x + 2Re[WD_{x,y}] + WD_y, \quad (2.5)$$

where $WD_{x,y}$ is the cross-Wigner distribution

$$WD_{x,y}(t,\omega) = \int_{-\infty}^{\infty} x(t+\frac{\tau}{2})y^*(t-\frac{\tau}{2})e^{-j\omega\tau}d\tau, \quad (2.6)$$

and $2Re[WD_{x,y}]$ is called a cross-term. Cross-terms can obscure other components or even closely-spaced auto-components, and they can also be mistaken for true signal components. In spite of the desirable properties of the Wigner distribution, the cross-terms make it so difficult to interpret that it is rarely used in applications.

The short-time Fourier transform and the Wigner distribution are related in a number of interesting ways. The squared magnitude of the short-time Fourier transform is the two-dimensional convolution of the Wigner distribution of the signal and the time-reversed complex conjugate of the analysis window [7].

$$|S(t,\omega)|^2 = WD_x ** WD_w^*(-t) \quad (2.7)$$

The squared magnitude of the short-time Fourier transform is thus a filtered version of the Wigner distribution in which the convolution smooths the time-frequency spectral estimate.

Furthermore, by making the change of variables $\tau = \frac{t}{2} + \frac{s}{2}$ in (2.1), we obtain

$$S_x(t, \omega) = \frac{e^{-j\frac{\omega t}{2}}}{2} \int_{-\infty}^{\infty} x\left(\frac{t}{2} + \frac{s}{2}\right) w\left(\frac{s}{2} - \frac{t}{2}\right) e^{-j\frac{\omega}{2}s} ds = \frac{e^{-j\frac{\omega t}{2}}}{2} WD_{x, w^{(-)}}\left(\frac{t}{2}, \frac{\omega}{2}\right) \quad (2.8)$$

and the short-time Fourier transform is shown to be merely a time- and frequency-scaled and phase-shifted cross-Wigner distribution. It is thus not surprising that when scaling effects are removed, the Wigner distribution and the short-time Fourier transform have the same effective signal resolution in time-frequency when certain windows are used [21].

2.6. Other Time-Frequency Representations

A large number of time-frequency representations have been introduced over the years in attempts to avoid some of the problems with the short-time Fourier transform and the Wigner distribution. Most of these representations have been shown to be members of Cohen's generalized class of time-frequency representations [8]. Other researchers have shown that most of these representations can be derived from the Wigner distribution. Few of these representations are considered today, because they either have worse performance than the short-time Fourier transform or they suffer from the same problems as the Wigner distribution but lack its nice properties [19]. The reader is referred to the dissertation by Boudreaux-Bartels [4] or the series of papers by Claasen and Mecklenbrauker [5-7] for a more complete survey of the various time-frequency representations.

The ambiguity function, defined as

$$AF_x(\tau, \phi) = \int_{-\infty}^{\infty} x\left(t + \frac{\tau}{2}\right) x^*\left(t - \frac{\tau}{2}\right) e^{-j\phi t} dt, \quad (2.9)$$

is of central importance in radar and detection theory. The ambiguity function is the two-dimensional Fourier transform of the Wigner distribution. As a bilinear representation, the

ambiguity function has cross-terms. Although it is closely related to the Wigner distribution, the ambiguity function lacks many of the desirable properties of the Wigner distribution, so it is rarely used in time-frequency analysis.

The only other time-frequency representation that attracts much interest today is the smoothed Wigner distribution [14], of which the pseudo-Wigner distribution and the squared magnitude of the short-time Fourier transform (see (2.7)) are special cases. Smoothed Wigner distributions are merely the Wigner distribution convolved with a two-dimensional filter $F(t,\omega)$ in time-frequency.

$$SWD_x(t,\omega) = WD_x(t,\omega) ** F(t,\omega) \quad (2.10)$$

The idea motivating this representation is that the smoothing function may filter out the cross-terms in the Wigner distribution while leaving the auto-components largely intact. The smoothed Wigner distribution allows a tradeoff between the high auto-component concentration of the Wigner distribution and the elimination of cross-terms in the short-time Fourier transform. As in the short-time Fourier transform, the choice of the filter function in the smoothed Wigner distribution has a major impact on the appearance and quality of the resulting time-frequency representation, and the best filter depends on the data. Since the effective resolution of the Wigner distribution and the short-time Fourier transform (with the proper window) are the same, the improvements offered by the smoothed Wigner distribution are largely cosmetic.

CHAPTER 3

A Data-Adaptive Time-Frequency Representation

As discussed in the previous chapter, the current time-frequency representations fall short of the ideal. The Wigner distribution has a number of desirable properties including high signal concentration in time-frequency, but the cross-terms make the Wigner distribution unsuitable for general use. The smoothed Wigner distribution allows flexibility in making tradeoffs between cross-terms and signal concentration, but it does not offer increased resolution over the Wigner distribution or the short-time Fourier transform with a properly chosen window.

The short-time Fourier transform performs very well in terms of concentration and resolution of a given signal component when a properly chosen window is used. However, the proper window function depends on the data, and no automated procedure currently exists for determining a good window. This means that the time-frequency analyst needs to know a lot about the signal before it can be processed effectively. A second problem is that for signals composed of several different components at different locations in time-frequency, the best window differs for each component. For a thorough analysis, several short-time Fourier transforms with different windows are needed. Furthermore, the analyst has no rational way to decide which window gives the most accurate representation of the data.

We propose a new representation that overcomes these problems with the short-time Fourier transform. The fact that different windows are needed for different signal components suggests the use of a data-adaptive window that varies at different time-frequency locations to achieve high concentration of any signal component at any location in time-

frequency. Since the Gaussian function is the most concentrated in time-frequency, we choose Gaussian windows. This leaves a two-parameter class (the real and imaginary part of the Gaussian window parameter) of functions that are equivalent in terms of their natural time-frequency concentration. They differ, though, in the time-frequency concentration they provide for a particular signal component. We define a measure of local signal concentration and compute the Gaussian window parameter maximizing this measure at each time-frequency location to achieve maximum concentration of the locally dominant signal component at every location in time-frequency. This procedure automates the choice of the window and thus overcomes the problem of window selection in the short-time Fourier transform.

The new time-frequency representation can be expressed as

$$A(t, \omega) = \int_{-\infty}^{\infty} x(\tau) \left[\frac{-Re[c_{t, \omega}]}{\pi} \right]^{\frac{1}{4}} e^{c_{i, \omega}(\tau-t)^2} e^{-j\omega\tau} d\tau, \quad (3.1)$$

Which is a projection of the signal $x(\tau)$ onto the Gaussian basis elements

$$\left[\frac{-Re[c_{t, \omega}]}{\pi} \right]^{\frac{1}{4}} e^{c_{i, \omega}(\tau-t)^2} e^{-j\omega\tau}. \quad (3.2)$$

The new representation is like a short-time Fourier transform with a Gaussian window, except that the Gaussian parameter $c_{t, \omega}$ may vary at different locations in time-frequency. It is important to note that, given $c_{t, \omega}$, the new representation is linear with respect to the signal, and thus it avoids the cross-terms associated with the Wigner distribution and other bilinear representations. Linearity also allows simple linear least squares techniques to be used to reconstruct the signal from the modified transform.

Although the new representation is a projection-based representation like the short-time Fourier transform and can be considered a composite of a number of short-time Fourier transforms, the variation of the Gaussian projection function in (3.2) with frequency means that the new representation does not produce the Fourier transform of the windowed signal. In situations in which the Gaussian projection function varies only with time, it seems reasonable to call the new representation an "adaptive window short-time Fourier transform." Even when the Gaussian function varies with frequency, it may be convenient to call this representation an adaptive window short-time Fourier transform, as long as it is understood that a Fourier transform is not actually present. In what follows, we use the terms "Gaussian projection function" and "window" interchangeably; the reader should note that the term "window" is not always strictly correct but represents an extension of the usual concept.

The performance of this new time-frequency representation depends on the selection of the adaptive Gaussian window parameters. In the following two chapters two different approaches to this problem are presented. In Chapter 4, we define a measure of local time-frequency signal concentration and find the Gaussian parameter maximizing this measure at each time-frequency location. This technique performs very well. Unfortunately, the optimal parameter is found by a brute-force search, so the computational requirements of this method are substantial. An alternate approach requiring much less computation is developed in Chapter 5. This method is based on the idea that a window matched to the data is known to yield high concentration of time-frequency concentrated components. Since Gaussian windows are used, we determine Gaussian parameters that most closely match those of the signal at each time-frequency location and apply these parameters in computing (3.2). This technique requires much less computation than the concentration measure approach, but its performance is inferior. The two methods allow a tradeoff between performance and computa-

tion time.

CHAPTER 4

Optimal Window Parameter Selection by Maximizing Concentration

4.1. Introduction

Maximum concentration of signal components in time-frequency is desired, since concentrated components in general overlap or interfere with other nearby components as little as possible and thus enhance resolution. Maximum concentration also implies that signals are confined as closely as possible to their proper support in time-frequency, which gives the interpreter more confidence in the time-frequency representation. Several measures of time-frequency concentration are available in the literature [15, 19, 32], but these measures are appropriate for describing the concentration in time-frequency of a signal consisting of a single component only. In this chapter, we define a new measure of concentration that appeals to the intuition for single-component signals and that can also be extended into a local measure of concentration, thus allowing the concentration of individual components in multi-component signals to be measured in the time-frequency regions they dominate. The Gaussian window parameter maximizing the local concentration is determined at each time-frequency location and applied in (3.1) to give a time-frequency representation with good concentration everywhere in time-frequency.

4.2. A Measure of Time-Frequency Signal Concentration

For a signal with a single concentrated component, the most concentrated representation is the least spread out in time-frequency. The time-frequency representation of a concentrated component tends to decay monotonically in all directions from a central peak. If the

representation is scaled to have a maximum value of unity, the most concentrated representation is then reasonably defined as the one with the minimum integral, or volume, of the normalized representation over the time-frequency plane. We suggest the following normalized volume V_N of the squared magnitude of the short-time Fourier transform $S(t,\omega)$ as a measure of time-frequency signal concentration, in which a smaller volume implies greater concentration.

$$V_N = \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |S(\tau,\phi)|^2 d\tau d\phi}{\max_{\tau,\phi} |S(\tau,\phi)|^2} \quad (4.1)$$

Since we are using Gaussian windows, we choose as the optimal Gaussian parameter c the one that satisfies

$$\min_c V_N = \min_c \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |S_c(\tau,\phi)|^2 d\tau d\phi}{\max_{\tau,\phi} |S_c(\tau,\phi)|^2}, \quad (4.2)$$

where S_c is the short-time Fourier transform using the Gaussian window with parameter c .

Having defined the measure of concentration, the definition of the optimal parameter in (4.2) is axiomatic. However, to have confidence in the reasonableness of the new measure of concentration as applied to this problem, we must at least confirm that it yields reasonable answers in certain special cases for which good answers are already known. It is well known that for Gaussian signals and windows that the matched window is optimal in terms of signal concentration according to the established concentration measures. As a partial confirmation of the reasonableness of the new measure, we demonstrate that the minimization in (4.2) yields the matched window when applied to Gaussian signals.

Consider the short-time Fourier transform of the Gaussian signal

$$x(t) = e^{bt^2}, \quad \text{Re}[b] = b_r < 0, \quad \text{Im}[b] = b_i = 0 \quad (4.3)$$

and the normalized Gaussian window

$$w(t) = \left[\frac{-2\text{Re}[c]}{\pi} \right]^{\frac{1}{4}} e^{ct^2}, \quad \text{Re}[c] < 0. \quad (4.4)$$

The squared magnitude of this short-time Fourier transform is divided by the maximum, which is located at the origin, to yield the normalized volume

$$V_N = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \exp \left[\frac{2(b_r c_r (b_r + c_r) + b_r c_i^2) t^2 - 2b_r c_i t \omega + \frac{1}{2}(b_r + c_r) \omega^2}{(b_r + c_r)^2 + c_i^2} \right] dt d\omega, \quad (4.5)$$

where $c_r = \text{Re}[c]$, $b_r = \text{Re}[b]$, and $c_i = \text{Im}[c]$. By completing the square in either t or ω , it can be shown that the integral

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \exp [At^2 + Bt\omega + C\omega^2] = \frac{2\pi}{\sqrt{4CA - B^2}}. \quad (4.6)$$

The normalized volume in (4.5) is thus

$$V_N = 2\pi \sqrt{\frac{(b_r + c_r)^2 + c_i^2}{b_r c_r}}, \quad (4.7)$$

from which it is immediately obvious that $c_i = 0$ is required for minimization. By taking the partial derivative with respect to c_r , it can be shown that

$$(c_r + b_r)(c_r - b_r) - c_i^2 = 0 \quad (4.8)$$

is required for a minimum. With $c_i = 0$, it is apparent that for $c_r = b_r$, (4.8) is satisfied. The reader can confirm by taking second derivatives that this is a minimizer. Thus the normalized volume is minimized for $c = b$, and the matched window yields the maximum concentration

as expected.

For simplicity, a Gaussian signal centered at the origin with no chirp component was used above. However, since the squared magnitude of the short-time Fourier transform of a time- and frequency-shifted Gaussian is accordingly shifted, and since a rotation in time-frequency of the signal and the window [19] merely rotates the magnitude of the short-time Fourier transform as well, no generality was lost in the preceding derivation, and the results hold for all Gaussians.

It should be noted at this point that the squared magnitude of the short-time Fourier transform is also called a cross-Ambiguity surface, and its volume over time-frequency is equal to the product of the energy of the signal and the window. The numerator in (4.2) is thus merely a constant for all values of c , and the parameter maximizing the maximum projection of the signal onto the window is the minimizer of the normalized volume. Thus the matched window is immediately seen to be the optimum. This simplification is not possible when the concentration measure is extended into a local concentration measure in the following sections, so for continuity this approach is avoided in the previous analysis.

4.3. A Measure of Local Time-Frequency Signal Concentration

The above scheme is appropriate for a single component and returns only one parameter for the entire signal. For signals with multiple components, a technique that locally selects the parameter to provide maximum concentration for the dominant component at each location in time-frequency is needed. We achieve this by defining a local measure of time-frequency signal concentration and then selecting the parameter maximizing this at each location in time-frequency.

The concentration measure in (4.2) is turned into a local concentration measure by multiplying the squared magnitude of the short-time Fourier transform by a localization weighting function

$$L_{c,k}(\tau-t, \phi-\omega) = \exp \left[kc_r(\tau-t)^2 + k \frac{(\phi-\omega-2c_i(\tau-t))^2}{4c_r} \right] \quad (4.9)$$

centered at (t, ω) . This function is largest at its center and decays monotonically in all radial directions from this point; the localization weighting function is in essence a two-dimensional window that causes only nearby components in time-frequency to influence the measure of concentration. This function is chosen because it is the Wigner distribution of the Gaussian window with parameter c (raised to the $\frac{k}{2}$ power), and it is thus a power of the time-frequency window leakage envelope for the Gaussian window [20]. This is important, because it weights the signal content at $(\tau-t, \phi-\omega)$ relative to the influence it has at (t, ω) , the location of interest. Since this influence varies with the parameter c , the localization weighting function must be a function of the window parameter c to avoid biasing the local concentration measure. The factor k determines the tightness of the localization weighting function; $k=0$ results in no localization, and a large k gives high localization but reduces the sensitivity of the concentration measure with respect to c . We generally choose k from 0.1 to 0.25.

After applying the localization weighting function to the squared magnitude of the short-time Fourier transform, we compute the normalized volume as in (4.1).

$$V_N(t, \omega) = \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} L_{c,k}(\tau-t, \phi-\omega) |S_c(\tau, \phi)|^2 d\tau d\phi}{\max_{\tau, \phi} L_{c,k}(\tau-t, \phi-\omega) |S_c(\tau, \phi)|^2} \quad (4.10)$$

The Gaussian parameters $c_{t, \omega}$ minimizing $V_N(t, \omega)$ are applied in (3.1) to yield the adaptive

time-frequency representation.

As before, we seek partial confirmation of the reasonableness of this local concentration measure by demonstrating that it yields the expected matched window when applied to a Gaussian signal. With rather extensive algebraic manipulation, the product in (4.10) of the squared magnitude of the short-time Fourier transform of the Gaussian signal in (4.3) and the localization window in (4.9) centered at (t, ω) can be placed in the form

$$|S(\tau, \phi)|^2 L_c(\tau-t, \phi-\omega) = N_{\max} \exp \left[A(\tau-\tau_0)^2 + B(\tau-\tau_0)(\phi-\phi_0) + C(\phi-\phi_0)^2 \right], \quad (4.11)$$

$$A = \frac{2b_r c_r [c_r(b_r+c_r)+c_i^2] + k[c_r^2+c_i^2][(b_r+c_r)^2+c_i^2]}{c_r[(b_r+c_r)+c_i^2]}, \quad (4.12)$$

$$B = \frac{-2b_r c_r c_i - \frac{1}{2} k c_i [(b_r+c_r)^2+c_i^2]}{c_r[(b_r+c_r)+c_i^2]}, \quad (4.13)$$

$$C = \frac{\frac{1}{2}(b_r+c_r)c_r + \frac{k}{4}[(b_r+c_r)^2+c_i^2]}{c_r[(b_r+c_r)+c_i^2]}, \quad (4.14)$$

where N_{\max} is the maximum

$$N_{\max} = \max_{\tau, \phi} |S(\tau, \phi)|^2 L_c(\tau-t, \phi-\omega), \quad (4.15)$$

which occurs at (τ_0, ϕ_0) . In computing the normalized volume in (4.10), the denominator is brought inside the integral to cancel N_{\max} , and the normalized volume can be found using (4.6). By taking first partial derivatives with respect to c_i and c_r , it can be shown that a critical point occurs at $c_r = b_r$, $c_i = 0$ for any value of k , the localization factor. That this point is a minimum is confirmed by showing that the Hessian matrix is positive definite. For the reasons discussed earlier for the global concentration measure, no generality is lost in this analysis, and thus minimizing the local concentration measure at any point in time-frequency

for any Gaussian yields the expected matched window parameters.

4.4. Implementation of the New Method

Although the concentration measures described in the previous two sections are defined and analyzed for continuous-time signals, these techniques are implemented on a digital computer with discrete-time signals. This is most easily accomplished with discrete approximations to the continuous-time operations; integrals are computed as summations, Gaussian functions are truncated, and the DFT is used in place of the Fourier transform.

As demonstrated earlier, it is possible with some difficulty to obtain analytically the optimal Gaussian parameter $c_{t,\omega}$ in certain special cases. Unfortunately, we have not found a general analytic expression for the optimal parameters for arbitrary signals. Furthermore, the function to be minimized is so extensive that use of a standard nonlinear minimization routine appears to be computationally prohibitive. We determine the optimal parameter using a brute force search over the entire range of reasonable parameters. Fortunately, the search grid for the Gaussian parameters c_r and c_i can be fairly coarse, since concentration is relatively insensitive to small errors in c .

For sampled signals, the parameter $c_{t,\omega}$ is determined by generating a discrete short-time Fourier transform and the localization weighting function on a time-frequency grid fine enough to yield a good approximation to the integral and the maximum in (4.10). The localization weighting function is applied to the short-time Fourier transform at each time-frequency location of interest, and the normalized volume is computed according to (4.10). This is done for all values of c to be searched, and the squared magnitude associated with the minimum normalized volume at each time-frequency location is the value of the new time-frequency representation at that location.

Since the computation involved is quite extensive, the necessary search ranges must be carefully considered. The optimal window parameter certainly need be computed only at time-frequency locations at which an output is desired. It is also possible to compute the optimal parameter on a grid coarser than the output grid and use the parameter at nearby locations. This is a reasonable thing to do if the optimal parameter is not expected to change much over the region on which it is used. The size of this region depends in general on the tightness of the localization weighting function, which is controlled by the factor k in (4.9). As a rough estimate, the spacing in samples between points at which the optimal parameter is determined should be about $k^{-1/2}$ in both the time and frequency directions.

The search grid for the optimal parameter c need not extend beyond the range of reasonable values. This range can generally be obtained by a preliminary glance at the time series or is known from the physics of the problem. For the approximation of the integral in (4.10) as a summation to be accurate, c_r should be such that the number of time samples for the window function to decay from the maximum to $\frac{1}{e}$ the maximum is not much less than three samples. The length of the DFT used in computing the short-time Fourier transform should be such that the same is true in the frequency direction as well, which limits the maximum value of c_r . The parameter c_i should be confined to a range such that no significant aliasing occurs in the short-time Fourier transform. This requirement is necessary to ensure that the integral in (4.10) is well approximated; aliasing chops off part of the signal component in the summation, which makes the partial volume appear less and can thus lead to an erroneous choice of the optimal parameter.

The search grid for c can be rather coarse, since signal concentration is relatively insensitive to small errors in c . The best choice of the search values depends on the application, but we use a grid that is linear in the square root of $-c_r$ from the maximum to minimum

search value, and on which the maximum and minimum c_i values are proportional to the square root of $-c_r$, with a linear grid over this range. It is hard to imagine an application in which more than 10 values of c_r and more than 20 values of c_i need be tried, for a maximum of 200 values to be searched. When available, additional knowledge about possible constraints can be incorporated to reduce the search. For example, if it is known that the signal components have no chirp, then c_i is set to zero, and only a one-dimensional search over c_r is required.

The grid on which the short-time Fourier transform is computed must be of sufficient density to ensure a reasonably accurate approximation to the integral in (4.10). The DFT need not be as long as the signal; the frequency sample density just needs to be sufficient. The user must also make sure that the boundaries of the computed short-time Fourier transform include the entire region in which the transform is essentially nonzero, again to ensure that the integral is accurately approximated.

The localization parameter k allows a tradeoff between the sensitivity of the concentration measure and the rate at which the Gaussian window parameter can vary. We have found that the concentration measure is generally not localized enough for $k < 0.1$, and too much sensitivity is sacrificed for $k > 0.5$. Within this range, a large k performs best for resolving the tops of nearby components, especially when they have different amplitudes and shapes. However, the lower levels of a component are often poorly concentrated when k is large. A small k should be used when the low-amplitude fringes of a component are of interest.

4.5. Extensions to Improve Computation Speed

Determination of the optimal Gaussian parameters at each time-frequency location requires a large amount of computation. The brute force search increases the computational

load by an amount proportional to the number of search values, but unfortunately this expense cannot be avoided. The computation of the normalized volume at each time-frequency location is by far the dominant computational load in the algorithm, and any means to speed this up would decrease the computational burden tremendously.

The numerator in the definition of the local normalized volume in (4.10) is a two-dimensional convolution of the localization window and the squared magnitude of the short-time Fourier transform. This convolution can be computed very efficiently by use of two-dimensional Fourier transforms, but unfortunately this technique cannot be used because the maximum of the product of the localization window and the short-time Fourier transform is also needed in the denominator of (4.10), so the product must be computed in the time-frequency domain to obtain this maximum.

To take advantage of efficient convolution, we modify the measure of concentration. The normalized volume can be thought of as the L_1 norm of the squared magnitude of the short-time Fourier transform multiplied by the localization window divided by the L_∞ norm of this function. This suggests that replacing the L_∞ norm in the denominator with another L_p norm may yield another function that effectively measures concentration but that is also easier to compute. Although any L_p norm, $p > 1$, works, we use the L_2 norm because it is easy to square the two functions. The modified volume V_M which we minimize is

$$V_M(t, \omega) = \frac{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} L_{c,k}(\tau-t, \phi-\omega) |S_c(\tau, \phi)|^2 d\tau d\phi}{\sqrt{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} L_{c,k}^2(\tau-t, \phi-\omega) |S_c(\tau, \phi)|^4 d\tau d\phi}}. \quad (4.16)$$

In practice, we minimize V_M^2 to avoid the square root. The denominator in (4.16) is now a two-dimensional convolution as well as the numerator, so fast convolution techniques can be

used to compute these functions.

As with the old concentration measure in (4.10), we confirm the reasonableness of this new measure by demonstrating that it yields the matched window parameters when applied to a Gaussian signal. This is easily shown by expressing (4.16) for the Gaussian signal and window from (4.3) and (4.4) in the form

$$V_M^2 = \frac{\left[\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} N_{\max} \exp[A(\tau-\tau_0)^2 + B(\tau-\tau_0)(\phi-\phi_0) + C(\phi-\phi_0)^2] d\tau d\phi \right]^2}{\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} N_{\max}^2 \exp[2A(\tau-\tau_0)^2 + 2B(\tau-\tau_0)(\phi-\phi_0) + 2C(\phi-\phi_0)^2] d\tau d\phi}, \quad (4.17)$$

where A , B , and C are the same as in (4.12), (4.13), and (4.14), respectively. From (4.6), we obtain the values of the integrals to get

$$V_M^2 = \frac{\left[\frac{N_{\max} \pi}{\sqrt{CA - \frac{B^2}{4}}} \right]^2}{\frac{N_{\max}^2 \pi}{\sqrt{4CA - \frac{(2B)^2}{4}}}} = \frac{2\pi}{\sqrt{CA - \frac{B^2}{4}}}, \quad (4.18)$$

which is the same function in (4.6) that is minimized by the matched window parameter in the original localized concentration measure. Thus the modified volume also yields the expected results everywhere.

The modified technique is implemented in the same way as described in Section 4.4, except that the time-frequency domain computation of the normalized volume at each location is replaced by two convolutions that compute all locations at once. The convolution arrays should be zero-padded sufficiently so that a linear rather than a cyclic convolution is obtained. The modified procedure speeds up the algorithm by one to two orders of magni-

tude for problems of a reasonable size. The new concentration measure generally performs slightly better than the original method as well, so we recommend its use even in a time-frequency domain implementation, although the difference is not substantial enough to override computational or convenience considerations.

CHAPTER 5

Optimal Window Selection by Gaussian Parameter Estimation

5.1. Introduction

Selection of the Gaussian window parameter by maximizing a local measure of concentration, as developed in the previous chapter, appeals to the intuition and performs quite well. However, the computational burden imposed by this technique is substantial, and situations may exist in which this burden is either unacceptable or in which some performance can be sacrificed in the interests of increased speed. In this chapter, we present an approach to selecting the Gaussian window parameter based on a Gaussian parameter estimate. This method requires much less computation than the method in Chapter 4; for situations in which the window varies only with time, the new method requires little more computation than the short-time Fourier transform with a fixed window. The performance of this technique, however, is not as good as the concentration-based method, and thus the two approaches offer a tradeoff between performance and computational burden.

In the Gaussian estimation method developed in this chapter, the goal of maximum signal concentration is approached indirectly. We know that the matched window gives optimal concentration of Gaussian signals, and we are using Gaussian windows in the new time-frequency representation in (3.1), so it seems reasonable that the Gaussian that most closely approximates a signal component will yield high concentration of that component in time-frequency. Thus we choose the window by finding the Gaussian that locally most closely approximates the signal. The performance and speed of such a technique depends largely on the Gaussian parameter estimation algorithm.

5.2. A Gaussian Parameter Estimate

Any direct least-squares local fit of the signal with a Gaussian function requires a non-linear minimization routine that imposes the computational burden we are trying to avoid. We therefore seek an alternate estimate of the Gaussian parameters that may not result in as good a fit but which is more attractive computationally. Such an estimate is analogous to linear predictive coding (LPC) in exponential fitting problems, which minimizes an "equation" error to determine the exponential parameters using simple linear minimization techniques.

We observe that for any discrete Gaussian signal $x(n) = e^{a+bn+cn^2}$, a , b , and c complex,

$$e^{2c} = \frac{x(n+1)x(n-1)}{x^2(n)}, \quad \text{all } n, \quad (5.1)$$

and the Gaussian parameter c can be determined from

$$c = \frac{1}{2} \ln \left[\frac{x(n+1)x(n-1)}{x^2(n)} \right], \quad \text{all } n. \quad (5.2)$$

This function serves as an estimate of the Gaussian parameters most closely matched to the signal at any time location n . In practice, we average several nearby estimates of e^{2c} to produce the window parameter that is actually used. We note in passing that this estimate is somewhat similar to instantaneous frequency. (We also note that a continuous-time non-linear differential equation analogous to this nonlinear difference equation is

$$2c = \frac{x''(t)}{x(t)} - \frac{[x'(t)]^2}{x^2(t)}, \quad (5.3)$$

for $x(t) = e^{a+bt+ct^2}$, a , b , and c complex.)

5.3. A Modified Gaussian Parameter Estimate With Frequency Selectivity

The Gaussian parameter estimate is accurate only for a single Gaussian component at any time location. For multi-component signals, a Gaussian bandpass filter $e^{-j\omega n + gn^2}$ is applied to retain only the dominant component at the current frequency of interest. The filtered signal

$$x_f(n) = x(n) * \left[\frac{2\text{Re}[g]}{\pi} \right]^{\frac{1}{4}} e^{gn^2} e^{-j\omega n} \quad (5.4)$$

is then analyzed as in (5.2) to obtain a Gaussian parameter $c_f(n)$ for the time-frequency location (n, ω) .

The Gaussian parameter obtained from (5.4) has been biased by the Gaussian filter, but for a Gaussian signal $x(n) = e^{a+bn+cn^2}$ the Gaussian parameter c_f estimated from $x_f(n)$ can be shown to be

$$c_f = \frac{cg}{c+g} \quad (5.5)$$

The unbiased parameter c can be calculated as

$$c = \frac{c_f g}{c_f - g} \quad (5.6)$$

and applied in (3.1) to give the time-frequency estimate at that point.

5.4. Implementation of the Parameter Estimation Method

If a single window parameter is desired at each time location, (5.2) is computed, and the resulting parameter is applied in (3.1) to yield the time-frequency representation. Several estimates of (5.1) from adjacent time locations may be averaged to reduce the variance of the Gaussian parameter estimate. This averaging introduces a tradeoff between rate of adaptation

and variance reduction in the Gaussian parameter estimate; the proper tradeoff depends on the situation and must be evaluated by the user. The algorithm is best implemented by computing and applying the best window at each successive time location. The window is the same for all frequencies at a given time, so an FFT can be used to compute the Fourier projections. Since little computation is needed to estimate the Gaussian parameter, this adaptive window short-time Fourier transform requires little more computation than a fixed window short-time Fourier transform.

If frequency selectivity is desired, the best implementation applies the Gaussian bandpass filter, computes the Gaussian window parameters for all times at that frequency, and computes the representation in (3.1) before moving on to the next frequency band. This implementation avoids two-dimensional arrays and allows the filtering to be done in the Fourier domain. Although a tight bandpass filter allows more frequency selectivity, it also increases the bias of the estimate and eventually degrades its performance. The Gaussian bandpass filter should be made as broad in frequency as is possible in the particular application. Again, averaging introduces a tradeoff that must be resolved by the user between the rate of adaptation and the variance of the Gaussian parameter estimates, but it should be noted that the Gaussian filtering already introduces smoothing of the filtered signal inversely proportional to the filter bandwidth, so less averaging is required for tighter bandpass filters. The computation required is on the order of that needed to compute a fixed window short-time Fourier transform using a DFT instead of an FFT; an FFT cannot be used since the window varies with frequency as well as with time.

An important component of a practical algorithm implementing this technique is a trap for unacceptable parameter estimates. For example, the Gaussian parameter estimate may yield a real part greater than zero. Our solution is to establish a range of acceptable values,

and any value outside the range is set to the closest acceptable value. Thus positive real parts of c are set to the maximum allowed (nonpositive) value for $Re[c]$.

CHAPTER 6

Applications and Comparisons

6.1. Introduction

This chapter describes applications of time-frequency analysis and examines the potential of the new adaptive techniques in these applications. Section 6.2 reviews major current applications of time-frequency analysis. The rest of the chapter describes applications of the adaptive time-frequency analysis techniques to the analysis of dispersive waves and compares their performance with the traditional techniques. In section 6.3, various time-frequency analysis techniques are applied to a synthetic atmospheric whistler. In section 6.4, the various time-frequency analysis techniques are applied to the analysis of a synthetic signal composed of two chirp Gaussians; this example is used by Cox and Mason in [9] to test their short-time maximum entropy spectral analysis technique. In the fifth section, a more sophisticated synthetic signal that was generated by Duckworth and Baggeroer from a wave equation model of the acoustic channel under Arctic ice is analyzed [11]. A real waveform recorded in the Arctic is analyzed in section 6.6 to test the performance and robustness of the adaptive technique on challenging real data.

6.2. Applications of Time-Frequency Analysis

Time-frequency analysis is used in a number of applications. Short-time spectral analysis plays a major, perhaps even dominant role in speech analysis, coding, and recognition [27]. Short-time Fourier and linear predictive analysis are the only time-frequency analysis techniques that are commonly used in speech applications. Although speech pro-

cessing is probably the most common application of time-frequency analysis, the time-frequency representation developed in this thesis is probably not useful in speech applications. The goal in speech time-frequency analysis is rarely high concentration of signal components; the window is either deliberately chosen to extensively smooth in the frequency direction to preserve only the formant structure, or to suppress the time direction pitch period effects in order to track the frequency variations in time-varying spectral features. Adaptive techniques could be developed that would somewhat enhance performance in the latter situation during formant frequency glides and other spectral transitions, but such techniques are unlikely to achieve broad usage. The goal of time-frequency analysis in speech applications is generally not to achieve high concentration, which is the goal of the adaptive techniques developed here, and thus these techniques are not very useful for speech applications.

Bioacoustics, or the study of animal sounds, is another application in which time-frequency analysis is often used. A variety of animals, including whales, dolphins, most bats, and many birds, produce frequency modulated or chirp sounds [22, 26, 29, 33]. Studies with dolphins indicate that they are able to discriminate frequency modulations better than any other waveform variations, which suggests that the frequency modulation is the most important feature of dolphin sounds in terms of communication [17]. As with dispersive modes, these waveforms are spread out in time and frequency but are concentrated in the time-frequency plane. The ability of the adaptive techniques to concentrate these waveforms and to track the varying frequency modulation in the waveforms makes these methods very applicable to bioacoustic time-frequency analysis.

Full waveform acoustic well logging is another application of much current interest for which time-frequency analysis techniques, and the adaptive method in particular, show great potential. The acoustic waveform is composed of several waveform components with differ-

ing characteristics, resulting from several different acoustic paths within the borehole complex. The Stonely wave component is usually dispersive. The goal of full waveform logging is to isolate and obtain information from all the components; current techniques generally determine only the time of the first arrival. Time-frequency analysis offers the possibility of doing this.

Time-frequency analysis is commonly used in the analysis of signals propagating in dispersive channels in which different frequencies either travel along different paths and/or with different velocities, causing different frequencies to be received at different times. Such signals are especially appropriate for time-frequency analysis, because they are spread out in both time and frequency but are concentrated in time-frequency. In general, individual non-dispersive transient signal components can be processed effectively using an appropriate combination of frequency filtering and time truncation, but time-frequency analysis is the only current technique offering high performance on dispersive signal components. Earthquake data is often dispersive, and time-frequency analysis techniques are widely used in these analyses [2, 10, 13]. Coal seams are usually dispersive acoustic channels, and time-frequency analysis techniques are often applied to coal seam data [9]. Techniques have been developed for removing dispersion from seismic waveforms [3], but these techniques require an accurate dispersion relation that is obtained by time-frequency analysis. The water channel under arctic ice is also a dispersive acoustic channel, and time-frequency analysis is used to determine the dispersion relations of the multiple acoustic modes [11]. The goal of time-frequency processing in these applications is to obtain enough resolution so that the various waveform components or modes can be distinguished, and to obtain the best possible concentration so that the dispersion relations can be obtained as accurately as possible. The adaptive techniques developed in this thesis have great potential in these applications.

6.3. Application to Synthetic Whistler Analysis

Whistling atmospherics are naturally occurring frequency modulated electromagnetic signals at acoustic frequencies. Storey studies these signals in some detail in [31] and concludes that they are electromagnetic waves induced by lightning strikes that propagate along a dispersive path through the ionosphere. This channel causes lower frequencies to arrive later in time, according to a dispersion law

$$t = Df^{-1/2}, \quad (6.1)$$

where D is the dispersion constant. Examination of his data reveals an amplitude dependency on time (or frequency) as well, but no theory for this dependency is presented in [31]. The real part of an analytic synthetic whistler

$$x(t) = \exp \left[-0.002(t-128)^2 + j \frac{2\pi 12^2}{t-64} \right] \quad (6.2)$$

with a Gaussian amplitude weight is shown in Figure 6.1. A contour plot of the Wigner distribution of this signal is shown in Figure 6.2a. The maximum value is normalized to 0dB, and the contours in this and subsequent plots are at 5dB intervals from -5dB to -25dB. The Wigner distribution is highly concentrated in time-frequency and accurately tracks the dispersion law. Inner cross-terms [18] between the "arms" of the waveform are visible in Figure 6.2a. For a relatively simple single-component signal such as this one, the cross-terms do not obscure the signal or confuse the interpreter, so the Wigner distribution is probably the best time-frequency representation to use.

Figure 6.2b contains a contour plot of the squared magnitude of the short-time Fourier transform of the synthetic whistler using the (non-chirped) Gaussian envelope of the signal in (6.2). The signal is poorly concentrated, and the dispersion relation is difficult to track accu-







Figure 6.1: Dispersive synthetic "whistler" $x(t) = \exp \left[-0.002(t-128)^2 + j \frac{2\pi 12^2}{t-64} \right]$




(a) Wigner Distribution



(b) Gaussian window STFT



(c) Matched window STFT



(d) Matched filter STFT

Figure 6.2: Wigner distribution and short-time Fourier transforms of a synthetic dispersive "whistler"

rately. Figure 6.2c shows the squared magnitude of the short-time Fourier transform using the matched window. The representation is quite concentrated, but inner cross-terms are visible in the plot. A more serious drawback is that the matched window representation does not track the dispersion relationship. The matched window representation thus may not be acceptable in some applications. Figure 6.2d displays the squared magnitude of the short-time Fourier transform using the matched filter. The central peak of this representation is quite sharp, which is expected because of the optimality of the matched filter for detection purposes. However, the matched filter imposes symmetry in the time-frequency representation on asymmetric signals such as this one; in this case, the signal appears to have an "S" shaped dispersion relation, which is completely wrong. For this reason, the matched filter is usually not acceptable in most time-frequency analysis applications.


Figure 6.3 contains time-frequency representations of the synthetic whistler with time and/or frequency varying filters. Figure 6.3a shows the result of using frequency varying Gaussian filters with the optimum bandwidth as derived by Barber and Ursell [1]. As pointed out by Storey, their analysis implies that the filters should have a one-sided bandwidth equal to

$$B(f) = \sqrt{\frac{df}{dt}}, \quad (6.3)$$


which yields a Gaussian filter parameter c in $h(t) = e^{ct^2}$

$$c = -\frac{f^{\frac{3}{2}}}{2D \ln 2}. \quad (6.4)$$


At the higher frequencies, this representation yields better resolution of the dispersion relation than the fixed Gaussian window in Figure 6.2b, but this analysis yields very poor results at the lower frequencies.




(a) Variable bandwidth (Barber)



(b) Parameter Estimation ATFR



(c) Maximum concentration ATFR



(d) Max. conc. without chirp

Figure 6.3: Adaptive time-frequency representations of a synthetic dispersive "whistler"

Figure 6.3b shows the result of using the adaptive window short-time Fourier transform using the Gaussian parameter estimation approach developed in Chapter 5. This representation accurately tracks the dispersion relation and has fairly good signal concentration. It also lacks the cross-terms in the Wigner distribution and the matched window short-time Fourier transform. Figure 6.3c shows the maximum concentration adaptive time-frequency representation developed in Chapter 4 with a localization weighting factor $k = 0.2$. This technique tracks the dispersion relation fairly well, but not quite as well as the Gaussian parameter estimation approach in regions where the dispersion rate is quickly varying. However, it yields a highly concentrated representation with roughly the concentration of the matched window representation that has no cross-terms. It also accurately tracks the dispersion relation. A larger localization weighting factor increases the ability of the maximum concentration approach to follow fast variations in the dispersion rate. The ability of the adaptive schemes to continuously track a nonlinear dispersion relation without introducing cross-terms or sacrificing concentration is a major advantage of the adaptive representations over the traditional time-frequency representations.

Figure 6.3d contains a contour plot of the concentration maximization adaptive time-frequency representation in which the Gaussian basis elements contain no chirp. The performance is much better than the fixed window short-time Fourier transforms in Figure 6.2, but not as good as the adaptive approaches in Figures 6.3b and 6.3c in which the chirp parameter varies. This illustrates the importance of the chirp component in obtaining high performance on dispersive signal components.

6.4. Analysis of Two Crossing Gaussian Chirps

Figure 6.4 contains the real part of a synthetic signal composed of two chirp Gaussian components with different parameters. This signal is almost identical to the synthetic signal




Figure 6.4: Signal composed of two Gaussian chirps

used by Cox and Mason to test a short-time maximum entropy time-frequency representation [9].

Figure 6.5 contains a contour plot of the squared magnitude of the short-time Fourier transform of this signal using a length-64 Hamming window, with the maximum value normalized to 0dB and contours at 5dB intervals from -25dB to -5dB. Unless otherwise noted, all the plots in this section are displayed in this manner. The tails of the two chirps begin to separate at about the -15dB level, but they are never very distinct. Figure 6.6 displays a short-time Fourier transform in which the window is matched to one of the chirp components. The matched component is very concentrated and distinct, but the other component is poorly concentrated. Furthermore, the apparent chirp rate of the other component is not the true chirp rate of that component; the chirp rate is biased by that of the window. In Figure 6.7 the window is matched to the other component. This component is now highly concentrated, but the first component is biased and poorly concentrated. A Gaussian window with parameters that are a compromise between the parameters of the two chirp components is used in Figure 6.8. The chirps are much better resolved than in the Hamming window short-time Fourier transform in Figure 6.5, but both signal components are biased and not extremely concentrated.

The Wigner distribution of this signal is shown in Figure 6.9. Both signal components are highly concentrated and distinct, but cross-terms appear between all the arms of the crossing chirps. In a simple synthetic example such as this, the cross-terms can be determined and ignored, but in more complicated situations they become a major problem. Figure 6.10 displays a short-time spectral estimate using Burg's technique (Maximum Entropy). The program developed by Marple in [23] was used. The contours are at 10dB intervals from -50dB to -10dB. The two chirps are fairly well resolved, and the chirp rates are not biased, but the




Figure 6.5: Hamming window short-time Fourier transform of signal composed of two chirp Gaussian components




Figure 6.6: Short-time Fourier transform of signal composed of two chirps with the window matched to one component




Figure 6.7: Short-time Fourier transform of signal composed of two chirps with the window matched to the other component




Figure 6.8: Short-time Fourier transform of signal composed of two chirps with another chirp
Gaussian window




Figure 6.9: Wigner distribution of signal composed of two chirps

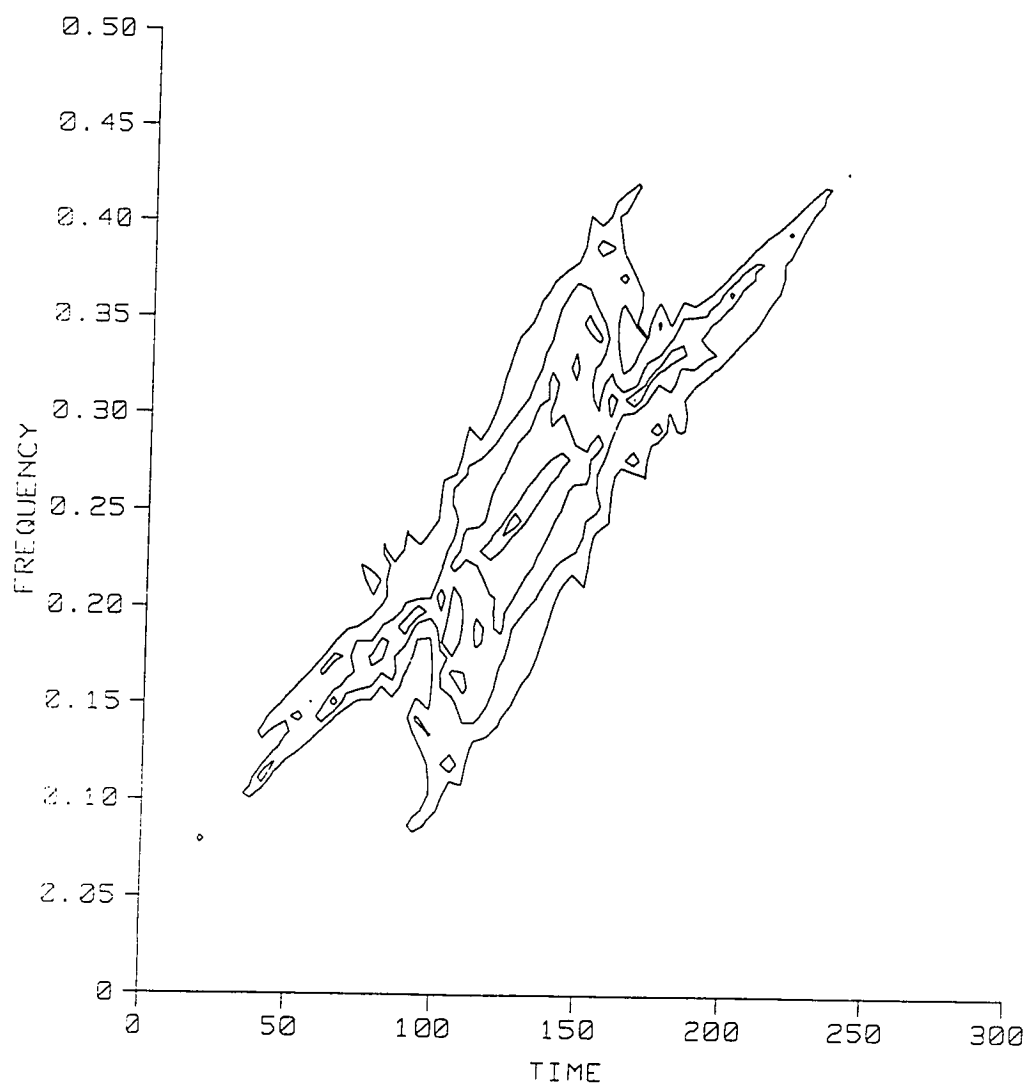


Figure 6.10: Maximum Entropy (Burg's Method) short-time spectral estimate of signal composed of two chirps

contours are very irregular. Figure 6.11 displays an adaptive time-frequency representation of the signal. The maximum concentration approach developed in Chapter 4 was used with a localization weighting factor of $k=0.1$ to obtain the Gaussian parameters. This method yields high concentration of both signal components in time-frequency. The chirp rates are unbiased, and the components are clearly separated. This representation is clearly superior to the fixed-window short-time Fourier transforms and the Maximum Entropy short-time spectral estimate in these respects, and it lacks the cross-terms that obscure the Wigner distribution.

6.5. Analysis of a Synthetic Arctic Acoustic Signal

Figure 6.12 contains a more realistic synthetic acoustic signal generated from a sophisticated model of the underwater acoustic channel under Arctic ice developed by Greg Duckworth and Art Baggeroer at MIT [11]. This signal consists of a dominant dispersive modal component between about time locations 100-200, a number of overlapping higher order dispersive modes at about times 30-100, and a nondispersive arrival at time locations 200-240. The main goal of time-frequency analysis of these signals is to determine the dispersion relations of the various modes. Figure 6.23, which is taken from Greg Duckworth's dissertation [11], plots the dispersion tracks on a short-time Fourier transform of a similar synthetic signal. In this section, we test the performance of the various time-frequency representations in this application.

The Wigner distribution of the synthetic "Arctic" signal is shown in Figure 6.13. As in previous sections, the contour levels in these figures are all at 5dB intervals from -25dB to -5dB. The dominant dispersive mode and the nondispersive late arrival are highly concentrated and are easily spotted in the Wigner distribution, but the higher order modes are completely obscured by cross terms between the first mode and the early portions of the signal.




Figure 6.11: Maximum concentration adaptive time-frequency representation of signal composed of two chirps




Figure 6.12: Synthetic acoustic signal propagating under Arctic ice ("Arctic signal")




Figure 6.13: Wigner distribution of Arctic signal

The Wigner distribution of a complicated signal is generally too complex to interpret.

The short-time Fourier transform of the signal using a length-64 Hamming window is shown in Figure 6.14. This window is not particularly appropriate for any of the signal components. The late arrival and the dominant dispersive mode are visible, but they are poorly concentrated. The other dispersive modes are jumbled together and cannot be distinguished in this representation. Figure 6.15 shows the short-time Fourier transform of the signal using a narrow-in-time Gaussian window. The late arrival and the dominant dispersive mode are very poorly concentrated by this technique. More detail can be seen in the region containing the less dispersive higher order modes, however; this window is more appropriate for this portion of the signal. Figure 6.16 contains the short-time Fourier transform using a chirp Gaussian window approximately matched to the large dispersive mode. This mode is highly concentrated in this representation, but the late arrival appears dispersive, and the higher modes are not distinct. A fundamental problem with these representations is that it is not clear which one offers the most accurate interpretation of the data.

Figure 6.17 contains an adaptive time-frequency representation of the synthetic arctic data using the concentration maximization technique with a localization weighting factor $k = 0.1$. The dominant dispersive mode is highly concentrated along its entire length, and the late arrival is fairly well concentrated, although the nearby dispersive mode has influenced the window parameter on the near side. A number of the higher order modes are visible in this representation; it clearly offers much better performance in this respect than any of the other representations. Figure 6.18 shows the adaptive representation using a localization weighting factor $k = 0.2$. It is quite similar to the previous plot, except that the lower contours are more irregular in some places, and the late arrival is less influenced by the large nearby dispersive component.




Figure 6.14: Short-time Fourier transform of Arctic signal with a length-64 Hamming window




Figure 6.15: Short-time Fourier transform of Arctic signal with a narrow-in-time Gaussian window




Figure 6.16: Short-time Fourier transform of Arctic signal with a chirp Gaussian window approximately matching the dominant modal component

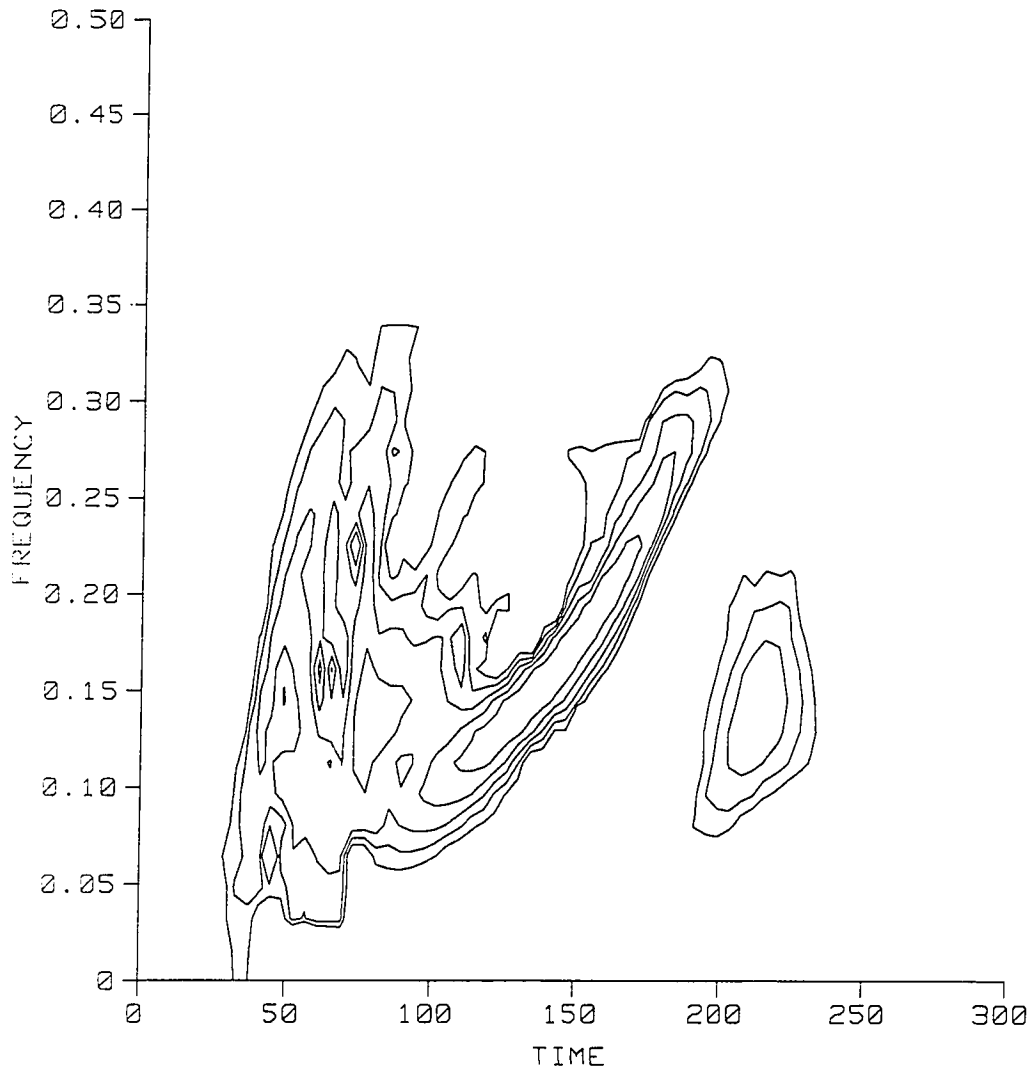


Figure 6.17: Maximum concentration adaptive time-frequency representation of Arctic signal with a localization weighting factor of 0.1




Figure 6.18: Maximum concentration adaptive time-frequency representation of Arctic signal with a localization weighting factor of 0.2

Figure 6.19 shows the adaptive time-frequency representation using the modified (L_1/L_2) measure of concentration developed in Chapter 4 with a localization weighting factor of $k=0.2$. This representation is extremely similar to the previous plot using the L_1/L_∞ concentration measure, which indicates that the modified measure is a reasonable measure of concentration even though it is less obvious conceptually than the original measure. Figure 6.20 shows the adaptive representation with the modified concentration measure computed using fast convolution in single precision. The convolution arrays were not zero-padded beyond the plot boundaries, which saves computation time; this accounts for the slight differences between this plot and the previous one. Figures 6.19 and 6.20 illustrate that the modified technique gives results practically identical to the original method, so in practice the fastest method should be used. On most machines, the modified method using fast convolution is fastest.

The final two figures in this section illustrate the performance of the Gaussian parameter estimation techniques developed in Chapter 5. Figure 6.21 shows the adaptive window short-time Fourier transform (time variation only of the window parameters) of the synthetic arctic signal. The Gaussian parameter estimates were smoothed with a length-15 triangular filter. The isolated dominant dispersive mode is well concentrated in this picture, and the latter portion of the late arrival is fairly clear, but the time-overlapping higher modes yield poor window parameter estimates that result in a poor representation of these regions. The Gaussian parameter estimation approach with frequency selectivity is shown in Figure 6.22. This method with frequency selectivity performs much better on signals with multiple components overlapping in time, but it does not perform as well as the maximum concentration approach. The amount of computation is greatly reduced, however, so a tradeoff exists between quality of the representation and computation time.




Figure 6.19: Maximum concentration adaptive time-frequency representation of Arctic signal with a localization weighting factor of 0.2, L_1/L_2 measure




Figure 6.20: Maximum concentration adaptive time-frequency representation of Arctic signal with a localization weighting factor of 0.2, fast convolution approach




Figure 6.21: Adaptive time-frequency representation of Arctic signal using the Gaussian parameter estimation approach; time variation only




Figure 6.22: Adaptive time-frequency representation of Arctic signal using the Gaussian parameter estimation approach with both time and frequency variation of window parameters

6.6. Analysis of a Real Arctic Acoustic Signal

In this section, we evaluate the performance of the adaptive technique on real data. We received this data from G.L. Duckworth and A.B. Baggeroer; it is a segment of an acoustic waveform propagating under the Arctic ice, and it results from a dynamite explosion 341 km away. This data was originally obtained as part of MIT and Woods Hole Oceanographic Institution participation in the FRAM II project supported by the Office of Naval Research. For consistency with the previous examples, the time and frequency axes are normalized. Figure 6.23 is a copy of a figure from Duckworth's dissertation [11], which shows a conventional short-time Fourier transform of a synthetic signal generated to match as closely as possible the real signal in Figure 6.24. The dispersion relations used in the model for the first seven modes are shown in Figure 6.23.

The short-time Fourier transform of the real arctic data using a length-64 Hamming window is shown in Figure 6.25. The dominant mode is easily visible, as is the high-frequency portion of the second mode. The high frequency tip of the third and perhaps the fourth mode are visible, but all other details are obscured. Figure 6.26 displays the adaptive time-frequency representation using the modified measure of concentration with a localization weighting factor of $k = 0.2$. Much more detail can be seen in this contour plot. More detail is visible in the dominant mode in Figure 6.26 than in the traditional short-time Fourier transform. The entire course of the second and third modes can be observed. The fourth and higher modes cross over each other and are too close to be completely resolved, but the center of the bundle of modes can be followed. This picture could be used to refine the estimates of the modes used to generate the synthetic signal.

This signal presents a real challenge to time-frequency analysis techniques. The background noise level is low, but the dispersive modes contain beats and other deviations from

- 72 -




Fig. 2.4) The spectrogram of the first trace of the data shown in figure 2.3a. A tapered window of .25 seconds length was used to achieve the time resolution using a simple discrete Fourier transform spectral estimator. The lines and characters on the plot are described in the text.

Figure 6.23: Copy of a figure from Duckworth's thesis showing the dispersion relations in a synthetic signal modelling the real data in Figure 6.24




Figure 6.24: Real arctic acoustic data from a shot 341 km from the receiver




Figure 6.25: Short-time Fourier transform of the real data with a length-64 Hamming window




Figure 6.26: Adaptive time-frequency representation of the real data

their synthetic counterparts. Isolated modes through almost overlapping modes cover a wide range of resolution difficulty. The fact that the adaptive techniques work well on this signal indicates that the technique is robust enough to work in challenging time-frequency analysis applications. This is essential, because the development of new time-frequency analysis techniques requiring more computation than the current schemes is justified only by increased performance on signals for which the current techniques are inadequate. Adaptive time-frequency representations offer that increased performance.

CHAPTER 7

Conclusion

7.1. Conclusions

The new adaptive time-frequency representations presented here obtain high resolution and concentration in time-frequency of all signal components. By adapting the window parameters, these methods overcome the limitations of short-time Fourier transforms with fixed windows. The new representations also avoid the cross-terms associated with the Wigner distribution and other bilinear representations. The method maximizing local concentration performs quite well on complicated, realistic test data and real data, and it is robust enough to work well in challenging practical applications with real data. This representation offers performance that often greatly exceeds that which can be obtained with any other time-frequency representation known to us. The only drawback to this technique is the large amount of computation it requires. The parameter estimation techniques work fairly well and require much less computation, so they may be useful in situations with relatively simple signals, in applications in which computing time is important, and when the absolute maximum performance is not required.

7.2. Suggestions for Further Research

The applications in Chapter 6 confirm the potential of adaptive time-frequency representations to achieve better performance than the current time-frequency representations. However, the particular methods developed in this dissertation may not be the best possible adaptive time-frequency representations. New methods for determining the Gaus-

sian parameters that represent different tradeoffs between performance and computation time would be useful. Of particular importance is the development of a high-performance concentration maximization approach that does not require as much computation as the current technique. It seems unlikely that an analytic solution for the Gaussian parameter maximizing the local concentration can be found for the measure of local concentration developed in this thesis. However, it is possible that other measures that capture the essence of local concentration yet better lend themselves to computation can be found. Efforts should be made in this direction.

This dissertation considers adaptive time-frequency representations based only on adaptive Gaussian windows as in (3.1). Non-projection-based approaches and methods based on other windows should be investigated. One wild idea is to do a local signal reconstruction at the time-frequency location of interest, and then apply this as the window. Other windows may lead to a decreased computational load or other advantages. The traditional adaptive filter literature may contain ideas that could be of use in adaptive time-frequency analysis.

The pseudo-Wigner distribution applies a two-dimensional filter to the Wigner distribution of a signal. As in the short-time Fourier transform, the quality of the pseudo-Wigner distribution depends on the proper choice of this filter, which depends on the data and may differ at different locations in time-frequency. Thus the development of adaptive filtering techniques for the pseudo-Wigner distribution, if successful, will result in the substantial improvements in performance found in the new adaptive techniques over the traditional short-time Fourier transform. This problem should certainly be investigated.

Application-specific simplifications and modifications of these techniques should be investigated. For example, the computation required by the concentration-based approach is reduced immensely if the signals are known to have no chirp component. Similarly, in

applications involving tracking frequency modulations of signals with a time extent greater than that feasible for the window length, adaptive techniques with a fixed window length and a variable chirp rate would be ideal. Similar simplifications can no doubt be made in many, if not most, of the applications in which these techniques may be useful. The objective of maximizing concentration appears to be appropriate for the multi-component dispersive wave analysis applications in Chapter 6, but in other applications, such as FM tracking, other optimization criteria may be more appropriate. This issue should be addressed in other applications.

References

- [1] N. F. Barber and F. Ursell, "The Response of a Resonant System to a Gliding Tone," *Phil. Mag.*, Vol.39, pp. 345-361, 1948.
- [2] S. Bloch and A. L. Hales, "New Techniques for the Determination of Surface Wave Phase Velocities," *Bull. Seism. Soc. Am.*, Vol.58, no.3, pp. 1021-1034, June 1968.
- [3] A. K. Booer, J. Chambers, and I. M. Mason, "A Fast Numerical Algorithm for the Recompression of Dispersed Seismic Signals," *Electronic Letters*, Vol.13, pp. 453-455, 1977.
- [4] G. F. Boudreaux-Bartels, "Time-Frequency Signal Processing Algorithms: Analysis and Synthesis Using Wigner Distributions," *Ph.D. dissertation, Rice University, Houston, TX*, December 1983.
- [5] T. A. C. M. Claasen and W. F. G. Mecklenbrauker, "The Wigner Distribution - A Tool for Time-Frequency Signal Analysis: Part I - Continuous-time Signals," *Philips J. Research*, no.3, pp. 217-250, 1980.
- [6] T. A. C. M. Claasen and W. F. G. Mecklenbrauker, "The Wigner Distribution - A Tool for Time-Frequency Signal Analysis: Part II - Discrete Time Signals," *Philips J. Research*, Vol.35, no.4/5, pp. 276-300, 1980.
- [7] T. A. C. M. Claasen and W. F. G. Mecklenbrauker, "The Wigner Distribution - A Tool for Time-Frequency Signal Analysis: Part III - Relations With Other Time-Frequency Signal Transformations," *Philips J. Research*, Vol.35, no.6, pp. 372-389, 1980.
- [8] L. Cohen, "Generalized Phase-Space Distribution Functions," *J. of Math. Phys.*, Vol.7, pp. 781-786, 1966.
- [9] K. B. Cox and I. M. Mason, "Maximum Entropy Analysis of Dispersed Seismic Waveforms," *Geophysics*, Vol.51, no.12, pp. 2225-2234, December 1986.
- [10] S. Crampin and M. Bath, "Higher Modes of Seismic Surface Waves: Mode Separation," *Geophysical Journal of the Royal Society*, Vol.10, pp. 81-92, 1965.
- [11] G. L. Duckworth, "Processing and Inversion of Arctic Ocean Refraction Data," *Massachusetts Institute of Technology Sc.D. dissertation*, 1983.
- [12] D. E. Dudgeon, "Detection of Narrowband Signals with Rapid Changes in Center Frequency," *IEEE ASSP 1984 Digital Signal Processing Workshop*, pp. 5.5.1-2, Chatham, MA, October 8-10, 1984.
- [13] A. Dziewonski, S. Bloch, and M. Landisman, "A Technique for the Analysis of Transient Seismic Signals," *Bull. Seismological Soc. Am.*, pp. 427-449, February 1969.

- [14] P. Flandrin, "Some Features of Time-Frequency Representations of Multicomponent Signals," *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing 1984*, pp. 41B.4.1-4, San Diego, March 19-21, 1984.
- [15] D. Gabor, "Theory of Communication," *Journal of the Institute of Electrical Engineers*, Vol.93, 1946.
- [16] C. W. Helstrom, "An Expansion of a Signal in Gaussian Elementary Signals," *IEEE Trans. Info. Theory*, Vol.IT-12, pp. 81-82.
- [17] L. M. Herman, "Cognition and Language Competencies of Bottlenosed Dolphins," in *Dolphin Cognition and Behavior: A Comparative Approach*, R. J. Schusterman, Eds. , pp. 221-252, 1986.
- [18] F. Hlawatsch, "Interference Terms in the Wigner Distribution," *International Conference of Digital Signal Processing*, Florence, Italy, September 5-8, 1984.
- [19] A. J. E. M. Janssen, "On the Locus and Spread of Pseudo-Density Functions in the Time-Frequency Plane," *Philips Journal of Research*, Vol.37, no.3, pp. 79-110, 1982.
- [20] D. L. Jones and T. W. Parks, "Time-Frequency Window Leakage in the Short-Time Fourier Transform," *To appear in Circuits, Systems, and Signal Processing*.
- [21] D. L. Jones and T. W. Parks, "Matched Windows in the Short-Time Fourier Transform," *IEEE ASSP 1986 Digital Signal Processing Workshop*, Chatham, Massachusetts, October 20-22, 1986.
- [22] D. E. Kroodsma, C. H. Miller, and H. Ouellet, *Acoustic Communication in Birds*, Academic Press , 1982.
- [23] S. L. Marple, *Digital Spectral Analysis with Applications*, Prentice-Hall, 1987.
- [24] S. R. McCaslin, "Spectral Estimation via Stationarity-Constrained Bandlimited Extrapolation," *Master's thesis, Rice University, Houston, TX*, October 1986.
- [25] J. Morlet, G. Arens, E. Fourgeau, and D. Giard, "Wave Propagation and Sampling Theory-Part II: Sampling Theory and Complex Waves," *Geophysics*, Vol.47, no.2, pp. 222-236, February 1982.
- [26] P. E. Purves and G. E. Pilleri, *Echolocation in Whales and Dolphins*, Academic Press, 1983.
- [27] L. R. Rabiner and R. W. Schafer, *Digital Processing of Speech Signals*. Englewood Cliffs, NJ: Prentice-Hall, 1978.
- [28] A. W. Rihaczek, *Principles of High-resolution Radar*. New York: McGraw Hill Book Co., 1969.
- [29] R. J. Schusterman, J. A. Thomas, and F. G. Wood, *Dolphin Cognition and Behavior: A Comparative Approach*, Lawrence Erlbaum Associates, Publishers, Hillsdale, NJ, 1986.
- [30] M. J. T. Smith and T. P. Barnwell, III, "A New Filter Bank Theory for Time-Frequency Representation," *IEEE Trans. Acoustics, Speech, and Signal Processing*, Vol.ASSP-35, no.3, pp. 314-327, March 1987.

- [31] L. R. O. Storey, "An Investigation of Whistling Atmospherics," *Philosophical Transactions of the Royal Society*, Vol.246, p. 113, 1953.
- [32] D. Ye. Vakman, "Optimum Signals Which Maximize Partial Volume Under an Ambiguity Surface," *Radio Engineering and Electronic Physics*, p. 1260, August 1967.
- [33] W. A. Wimsalt, *Biology of Bats: Volume II*, Academic Press, 1970.
- [34] N. G. deBruijn, "A Theory of Generalized Functions with Applications to Wigner Distribution and Weyl Correspondence," *Nieuw Archief voor Wiskunde*, Vol.21, pp. 205-208, 1973.