


Time-Frequency Toolbox

For Use with MATLAB



*François Auger **

*Patrick Flandrin **

Paulo Gonçalves °

*Olivier Lemoine **

** CNRS (France)*

° Rice University (USA)

1995-1996



Copyright (C) 1996 CNRS (France) and Rice University (USA).
Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

The Time-Frequency Toolbox has been mainly developed under the auspices of the French CNRS (*Centre National de la Recherche Scientifique*). It results from a research effort conducted within its Groupements de Recherche "Traitement du Signal et Images" (O. Macchi) and "Information, Signal et Images" (J.-M. Chassery). Parts of the Toolbox have also been developed at Rice University, when one of the authors (PG) was visiting the Department of Electrical and Computer Engineering, supported by NSF. Supporting institutions are gratefully acknowledged, as well as M. Guglielmi, M. Najim, R. Settineri, R.G. Baraniuk, M. Chausse, D. Roche, E. Chassande-Mottin, O. Michel and P. Abry for their help at different phases of the development.

Contents

1	Introduction	9
1.1	Presentation	9
1.2	Background, system requirements and installation	10
1.3	Introductory examples	10
1.3.1	Example 1	10
1.3.2	Example 2	13
1.3.3	Example 3	14
2	Non stationary signals	19
2.1	Time representation and frequency representation	19
2.2	Localization and the Heisenberg-Gabor principle	20
2.2.1	Example 1	20
2.2.2	Example 2	21
2.3	Instantaneous frequency	22
2.4	Group delay	24
2.5	About stationarity	25
2.6	How to synthesize a mono-component non-stationary signal . .	26
2.7	What about multi-component non-stationary signals ?	29
3	First class of solutions : the atomic decompositions	33
3.1	The Short-Time Fourier Transform	33
3.1.1	Definition	33
3.1.2	An example	35
3.1.3	Some properties	36
3.1.4	Time-frequency resolution	37
3.2	Time-scale analysis and the wavelet transform	40
3.2.1	Definitions and interpretation	41
3.2.2	Properties	42
3.3	Sampling considerations	43

3.3.1	The discrete STFT	43
3.3.2	The Gabor Representation	44
3.3.3	The discrete wavelet transform	46
3.4	From atomic decompositions to energy distributions	48
3.4.1	The spectrogram	48
3.4.2	The scalogram	52
3.4.3	Conclusion	54
4	Second class of solutions : the energy distributions	57
4.1	The Cohen's class	58
4.1.1	The Wigner-Ville distribution	58
4.1.2	The Cohen's class	67
4.1.3	Link with the narrow-band ambiguity function	72
4.1.4	Other important energy distributions	76
4.1.5	Conclusion	82
4.2	The affine class	83
4.2.1	Axiomatic definition	83
4.2.2	Some examples	86
4.2.3	Relation with the ambiguity domain	95
4.2.4	The affine Wigner distributions	98
4.2.5	The pseudo affine Wigner distributions	102
4.2.6	Conclusion	107
4.3	The reassignment method	108
4.3.1	Introduction	108
4.3.2	The reassignment of the spectrogram	109
4.3.3	Reassignment of the Cohen's class representations	111
4.3.4	Reassignment of the affine class representations	113
4.3.5	Numerical examples	113
4.3.6	Connected approaches	115
4.3.7	Conclusion	116
5	Extraction of information from a time-frequency image	123
5.1	Moments and marginals	123
5.1.1	Moments	123
5.1.2	Marginals	124
5.2	More on interferences: information on phase	124
5.3	Renyi information	126
5.4	Time-frequency analysis: help to decision	128
5.4.1	General considerations	128

5.4.2	An example: detection and estimation of linear FM signals	129
5.5	Analysis of local singularities	132
GNU Free Documentation License		138
1.	Applicability and definitions	139
2.	Verbatim copying	141
3.	Copying in quantity	142
4.	Modifications	142
5.	Combining documents	144
6.	Collections of documents	145
7.	Aggregation with independent works	145
8.	Translation	146
9.	Termination	146
10.	Future revisions of this license	146
	ADDENDUM: How to use this License for your documents	147
Bibliography		148

Chapter 1

Introduction

1.1 Presentation

The Time-Frequency Toolbox is a collection of M-files developed for the analysis of non-stationary signals using time-frequency distributions. This toolbox includes two groups of files :

- the signal generation files, which allow the synthesis of numerous kinds of non-stationary signals ;
- the processing files, including the time-frequency distributions and other related processing functions.

As usual under MATLAB, each function of the toolbox has a help entry that you can refer to by typing

```
>> help name_of_the_file
```

at the prompt of the matlab command window. In almost every case, a simple example is given, which facilitates the use of the function.

Seven demonstration M-files are also available, which provide sequences of examples illustrating the possibilities of the Time-Frequency Toolbox, and following closely the plan of this tutorial. These files are :

tfdemo	Main menu of the demonstration
tfdemo1	Introduction
tfdemo2	Non-stationary signals
tfdemo3	Linear time-frequency representations
tfdemo4	Cohen's class time-frequency distributions
tfdemo5	Affine class time-frequency distributions
tfdemo6	Reassigned time-frequency distributions
tfdemo7	Extraction of information

The aim of this Tutorial is to present the way to use the Time-Frequency Toolbox, and also to introduce the reader in an illustrative and friendly way to the theory of time-frequency analysis. We advise the reader, when looking at a chapter of this tutorial, to run simultaneously the corresponding demonstration file. In this way, he will have a good understanding of the Toolbox.

1.2 Background, system requirements and installation

This Toolbox is primarily intended for researchers and engineers with some knowledge on signal processing theory. In particular, the concepts of Fourier transform, Shannon sampling and stationarity are important to understand the following features.

The Time-Frequency Toolbox assumes that MATLAB v.4.2c (or a later version) is present on your system, as well as the Signal Processing Toolbox v.3.0 (or a later version).

Instructions for installing this toolbox on a workstation or a large machine are found in the MATLAB Installation Guide. Instructions for installing on micro computers are found in the MATLAB User's Guide.

1.3 Introductory examples

1.3.1 Example 1

Let us consider first a signal with constant amplitude, and with a linear frequency modulation varying from 0 to 0.5 in normalized frequency (ratio of the frequency in Hertz to the sampling frequency, with respect to the Shannon sampling theorem). This signal is called a chirp, and as its frequency content is varying with time, it is a non-stationary signal. To obtain such a signal, we can use the M-file `fmlin.m`, which generates a linear frequency modulation (see fig. 1.1):

```
>> sig1=fmlin(128,0,0.5);  
>> plot(real(sig1));
```

From this time-domain representation, it is difficult (except for experienced specialists) to say what kind of modulation is contained in this signal : what are the initial and final frequencies, is it a linear, parabolic, hyperbolic... frequency modulation ?




Figure 1.1: Linear frequency modulation (chirp)

If we now consider the energy spectrum of this signal `sig1` by squaring the modulus of its Fourier transform (using the `fft` function) (see fig. 1.2),

```
>> dsp1=fftshift(abs(fft(sig1)).^2);
>> plot((-64:63)/128,dsp1);
```




Figure 1.2: Energy spectrum of the chirp

we still can not say, from this plot, anything about the evolution in time of the frequency content. This is due to the fact that the Fourier transform is a decomposition on complex exponentials, which are of infinite duration and completely unlocalized in time. Time information is in fact encoded in the phase of the Fourier transform (which is simply ignored by the energy

spectrum), but their interpretation is not straightforward and their direct extraction is faced with a number of difficulties such as phase unwrapping. In order to have a more informative description of such signals, it would be better to directly represent their frequency content while still keeping the time description parameter : this is precisely the aim of time-frequency analysis. To illustrate this, let us try the Wigner-Ville distribution on this signal (see fig. 1.3):

```
>> tfrwv(sig1);
```




Figure 1.3: Wigner-Ville distribution of the chirp

Without going into details about this representation (it will be developed in the following), we can see that the linear progression of the frequency with time, from 0 to 0.5, is clearly shown.

If we now add some complex white gaussian noise on this signal, using the M-files `noisecg.m` and `sigmerge.m`, with a 0 dB signal to noise ratio (see fig. 1.4),

```
>> sig2=sigmerge(sig1,noisecg(128),0);
>> plot(real(sig2));
```

and consider the spectrum of it (see fig. 1.5):

```
>> dsp2=fftshift(abs(fft(sig2)).^2);
>> plot((-64:63)/128,dsp2);
```




Figure 1.4: Chirp embedded in a 0 dB white gaussian noise

it is worse than before to interpret these plots. On the other hand, the Wigner-Ville distribution still show quite clearly the linear progression of the frequency with time (see fig. 1.6) :

```
>> tfrwv(sig2);
```

1.3.2 Example 2

The second example we consider is a bat sonar signal, recorded with a sampling frequency of 230.4 kHz and an effective bandwidth of [8 kHz, 80 kHz] (this recording was part of the research program RCP 445 supported by CNRS (Centre National de la Recherche Scientifique, France) [Fla86]).

First, load the signal from the MAT-file `bat.mat` (see fig. 1.7) :

```
>> load bat
>> t0=linspace(0,2500/2304,2500);
>> plot(t0,bat); xlabel('Time [ms]');
```

From this plot, we can not say precisely what is the frequency content at each time instant t ; similarly, if we look at its spectrum (see fig. 1.8),

```
>> dsp=fftshift(abs(fft(bat)).^2);
>> f0=(-1250:1249)*230.4/2500;
>> plot(f0,dsp); xlabel('Frequency [kHz]');
```

we can not say at what time the signal is located around 38 kHz, and at what time around 40 kHz (you can use the zoom function to see more precisely what




Figure 1.5: Energy spectrum of the noisy chirp

is happening around these frequencies ; see the Matlab Reference Guide). Let us now consider a representation called the pseudo Wigner-Ville distribution, applied on the most interesting part of this signal (this distribution was obtained with the M-file `tfrpwv.m`, stored in the matrix `tfr` and saved with the signal in the MAT-file `bat.mat` ; the corresponding time- and frequency-samples `t` and `f` where also saved on `bat.mat`) (see fig. 1.9) :

```
>> contour(t,f,tfr,5); axis('xy');
>> xlabel('Time [ms]'); ylabel('Frequency [kHz]');
>> title('TFRPWV of a bat signal');
```

We then have a nice description of its spectral content varying with time : it is a narrow-band signal, whose frequency content is decreasing from around 55 kHz to 38 kHz, with a non-linear frequency modulation (approximately of hyperbolic shape).

1.3.3 Example 3

The last introductory example presented here is a transient signal embedded in a -5 dB white gaussian noise. This transient signal is a constant frequency modulated by a one-sided exponential amplitude (see fig. 1.10) :

```
>> trans=amexpols(64).*fmconst(64);
>> sig=zeros(100,1) ; trans ; zeros(92,1)];
>> sign=sigmerge(sig,noiseCG(256),-5);
>> plot(real(sign));
>> dsp=fftshift(abs(fft(sign)).^2);
>> plot((-128:127)/256,dsp);
```




Figure 1.6: Wigner-Ville distribution of the noisy chirp

From these representations, it is difficult to localize precisely the signal in the time-domain as well as in the frequency domain. Now let us have a look at the spectrogram of this signal calculated using the M-file `tfrsp.m` (see fig. 1.11):

```
>> tfrsp(sign);
```

the transient signal appears distinctly around the normalized frequency 0.25, and between time points 125 and 160.




Figure 1.7: Sonar signal from a bat




Figure 1.8: Energy spectrum of the bat sonar signal




Figure 1.9: Pseudo-WVD of the bat sonar signal




Figure 1.10: Time- and frequency- representation of a noisy transient signal




Figure 1.11: Spectrogram of the noisy transient signal

Chapter 2

Non stationary signals

This chapter presents some useful definitions that constitute the background of time-frequency analysis (most of the information presented in this tutorial are extracted from [Fla93]). After a brief recall on time-domain and frequency-domain representations, we introduce the concepts of time and frequency localizations, time-bandwidth product and the constraint associated to this product (the Heisenberg-Gabor inequality). Then, the instantaneous frequency and the group delay are presented as a first solution to the problem of time localization of the spectrum. We carry on by defining non-stationarity from its opposite, stationarity, and show how to synthesize such non-stationary signals with the toolbox. Finally, we show that in the case of multi-component signals, these mono-dimensional functions (instantaneous frequency and group delay) are not sufficient to represent these signals ; a two-dimensional description (function of time *and* frequency) is necessary.

2.1 Time representation and frequency representation

The time representation is usually the first (and the most natural) description of a signal we consider, since almost all physical signals are obtained by receivers recording variations with time.

The frequency representation, obtained by the *Fourier transform*

$$X(\nu) = \int_{-\infty}^{+\infty} x(t) e^{-j2\pi\nu t} dt,$$

is also a very powerful way to describe a signal, mainly because the relevance of the concept of frequency is shared by many domains (physics, astronomy, economics, biology ...) in which periodic events occur.

But if we look more carefully at the spectrum $X(\nu)$, it can be viewed as the coefficient function obtained by expanding the signal $x(t)$ into the family of infinite waves, $\exp\{j2\pi\nu t\}$, which are completely unlocalized in time. Thus, the spectrum essentially tells us which frequencies are contained in the signal, as well as their corresponding amplitudes and phases, but does not tell us at which times these frequencies occur.

2.2 Localization and the Heisenberg-Gabor principle

A simple way to characterize a signal simultaneously in time and in frequency is to consider its mean localizations and dispersions in each of these representations. This can be obtained by considering $|x(t)|^2$ and $|X(\nu)|^2$ as probability distributions, and looking at their mean values and standard deviations :

$$\begin{aligned} t_m &= \frac{1}{E_x} \int_{-\infty}^{+\infty} t |x(t)|^2 dt && \text{average time} \\ \nu_m &= \frac{1}{E_x} \int_{-\infty}^{+\infty} \nu |X(\nu)|^2 d\nu && \text{average frequency} \\ T^2 &= \frac{4\pi}{E_x} \int_{-\infty}^{+\infty} (t - t_m)^2 |x(t)|^2 dt && \text{time spreading} \\ B^2 &= \frac{4\pi}{E_x} \int_{-\infty}^{+\infty} (\nu - \nu_m)^2 |X(\nu)|^2 d\nu && \text{frequency spreading} \end{aligned}$$

where E_x is the *energy* of the signal, assumed to be finite (bounded) :

$$E_x = \int_{-\infty}^{+\infty} |x(t)|^2 dt < +\infty.$$

Then a signal can be characterized in the time-frequency plane by its mean position (t_m, ν_m) and a domain of main energy localization whose area is proportional to the *time-bandwidth product* $T \times B$.

2.2.1 Example 1

These time and frequency localizations can be evaluated thanks to the M-files `loctime.m` and `locfreq.m` of the Toolbox. The first one gives the average time center (t_m) and the duration (T) of a signal, and the second one the average normalized frequency (ν_m) and the normalized bandwidth (B). For example, for a linear chirp with a gaussian amplitude modulation, we obtain (see fig. 2.1) :

```
>> sig=fmlin(256).*amgauss(256);
>> [tm,T]=loctime(sig)           --->  tm=128      T=32
>> [num,B]=locfreq(sig)          --->  num=0.249   B=0.0701
```




Figure 2.1: Linear chirp with a gaussian amplitude modulation

One interesting property of this product $T \times B$ is that it is lower bounded

:

$$T \times B \geq 1.$$

This constraint, known as the *Heisenberg-Gabor inequality*, illustrates the fact that a signal can not have simultaneously an arbitrarily small support in time and in frequency. This property is a consequence of the definition of the Fourier transform. The lower bound $T \times B = 1$ is reached for gaussian functions :

$$x(t) = C \exp[-\alpha(t - t_m)^2 + j2\pi\nu_m(t - t_m)]$$

with $C \in \mathcal{R}$, $\alpha \in \mathcal{R}_+$. Therefore, the gaussian signals are those which minimize the time-bandwidth product according to the Heisenberg-Gabor inequality.

2.2.2 Example 2

To check the Heisenberg-Gabor inequality numerically, we consider a gaussian signal and calculate its time-bandwidth product (see fig. 2.2) :

```
>> sig=amgauss(256);
>> [tm,T]=loctime(sig);
```

```
>> [fm,B]=locfreq(sig);
>> [T,B,T*B]          ---->  T=32  B=0.0312  T*B=1
```




Figure 2.2: gaussian signal : lower bound of the Heisenberg-Gabor inequality

Hence, the time-bandwidth product obtained, when using the file `amgauss.m`, is minimum.

2.3 Instantaneous frequency

Another way to describe a signal simultaneously in time and in frequency is to consider its *instantaneous frequency*. In order to introduce such a function, we must define first the concept of *analytic signal*.

For any real valued signal $x(t)$, we associate a complex valued signal $x_a(t)$ defined as

$$x_a(t) = x(t) + jHT(x(t))$$

where $HT(x)$ is the *Hilbert transform* of x (x_a can be obtained using the M-file `hilbert.m` of the Signal Processing Toolbox). $x_a(t)$ is called the analytic signal associated to $x(t)$. This definition has a simple interpretation in the frequency domain since X_a is a single-sided Fourier transform where the negative frequency values have been removed, the strictly positive ones have been doubled, and the DC component is kept unchanged :

$$X_a(\nu) = 0 \quad \text{if } \nu < 0$$

$$\begin{aligned}
 X_a(\nu) &= X(0) & \text{if } \nu = 0 \\
 X_a(\nu) &= 2X(\nu) & \text{if } \nu > 0
 \end{aligned}$$

(X is the Fourier transform of x , and X_a the Fourier transform of x_a). Thus, the analytic signal can be obtained from the real signal by forcing to zero its spectrum for the negative frequencies, which do not alter the information content since for a real signal, $X(-\nu) = X^*(\nu)$.

From this signal, it is then possible to define in a unique way the concepts of *instantaneous amplitude* and *instantaneous frequency* by :

$$\begin{aligned}
 a(t) &= |x_a(t)| && \text{instantaneous amplitude} \\
 f(t) &= \frac{1}{2\pi} \frac{d \arg x_a(t)}{dt} && \text{instantaneous frequency}
 \end{aligned}$$

An estimation of the instantaneous frequency is given by the M-file `instfreq.m` of the Time-Frequency toolbox :

Example (see fig. 2.3)

```

>> sig=fmlin(256); t=(3:256);
>> ifr=instfreq(sig); plotifl(t,ifr);

```




Figure 2.3: Estimation of the instantaneous frequency of a linear chirp

As we can see from this plot, the instantaneous frequency shows with success the evolution with time of the frequency content of this signal.

2.4 Group delay

The instantaneous frequency characterizes a local frequency behavior as a function of time. In a dual way, the local time behavior as a function of frequency is described by the *group delay* :

$$t_x(\nu) = -\frac{1}{2\pi} \frac{d \arg X_a(\nu)}{d\nu}.$$

This quantity measures the average time arrival of the frequency ν . The M-file `sgrpdlay.m` of the Time-Frequency Toolbox gives an estimation of the group delay of a signal (do not mistake it for the file `grpdelay.m` of the signal processing toolbox which gives the group delay of a digital filter). For example, with signal `sig` of the previous example, we obtain (see fig. 2.4) :

```
>> sig=fmlin(256); fnorm=0:.05:.5;
>> gd=sgrpdlay(sig,fnorm); plot(gd,fnorm);
```




Figure 2.4: Estimation of the group delay of the previous chirp

Be careful of the fact that in general, instantaneous frequency and group delay define two different curves in the time-frequency plane. They are approximately identical only when the time-bandwidth product $T \times B$ is large. To illustrate this point, let us consider a simple example. We calculate the instantaneous frequency and group delay of two signals, the first one having a large $T \times B$ product, and the second one a small $T \times B$ product (see fig. 2.5) :

```
>> t=2:255;
```

```

>> sig1=amgauss(256,128,90).*fmlin(256,0,0.5);
>> [tm,T1]=loctime(sig1); [fm,B1]=locfreq(sig1);
>> T1*B1          --->  T1*B1=15.9138
>> ifr1=instfreq(sig1,t); f1=linspace(0,0.5-1/256,256);
>> gd1=sgrpdlay(sig1,f1); plot(t,ifr1,'*',gd1,f1,'-')
>> sig2=amgauss(256,128,30).*fmlin(256,0.2,0.4);
>> [tm,T2]=loctime(sig2); [fm,B2]=locfreq(sig2);
>> T2*B2          --->  T2*B2=1.224
>> ifr2=instfreq(sig2,t); f2=linspace(0.2,0.4,256);
>> gd2=sgrpdlay(sig2,f2); plot(t,ifr2,'*',gd2,f2,'-')

```




Figure 2.5: Estimation of the instantaneous frequency (stars) and group delay (line) of two different chirps with different amplitude modulations. The first plot corresponds to a large $T \times B$ product while the second corresponds to a small one

On the first plot, the two curves are almost superimposed (i.e. the instantaneous frequency is the inverse transform of the group delay), whereas on the second plot, the two curves are clearly different.

2.5 About stationarity

Before talking about non-stationarity, which is a 'non-property', we must define what we call *stationarity*.

A deterministic signal is said to be *stationary* if it can be written as a discrete sum of sinusoids :

$$x(t) = \sum_{k \in \mathcal{N}} A_k \cos [2\pi\nu_k t + \Phi_k] \quad \text{for a real signal}$$

$$x(t) = \sum_{k \in \mathcal{N}} A_k \exp [j(2\pi\nu_k t + \Phi_k)] \quad \text{for a complex signal}$$

i.e. as a sum of elements which have constant instantaneous amplitude and instantaneous frequency.

In the random case, a signal $x(t)$ is said to be *wide-sense stationary* (or stationary up to the second order) if its expectation is independent of time and its autocorrelation function $E[x(t_1)x^*(t_2)]$ depends only on the time difference $t_2 - t_1$. We can then show that the associated analytic signal has constant instantaneous amplitude and frequency expectations, which can be connected to the deterministic case.

So a signal is said to be *non-stationary* if one of these fundamental assumptions is no longer valid. For example, a finite duration signal, and in particular a *transient signal* (for which the length is short compared to the observation duration), is non-stationary.

2.6 How to synthesize a mono-component non-stationary signal

One part of the Time-Frequency Toolbox is dedicated to the generation of non-stationary signals. In that part, three groups of M-files are available :

1. The first one allows to synthesize different amplitude modulations. These M-files begin with the prefix 'am'. For example, `amrect.m` computes a rectangular amplitude modulation, `amgauss.m` a gaussian amplitude modulation ...
2. The second one proposes different frequency modulations. These M-files begin with 'fm'. For example, `fmconst.m` is a constant frequency modulation, `fmhyp.m` a hyperbolic frequency modulation ...
3. The third one is a set of pre-defined signals. Some of them begin with 'ana' because these signals are analytic (for example `anastep`, `anabpsk`, `anasing` ...), other have special names (`doppler`, `atoms` ...).

The first two groups of files can be combined to produce a large class of non-stationary signals, multiplying an amplitude modulation and a frequency modulation.

Examples

We can multiply the linear frequency modulation of Example 1 (see page 20) by a gaussian amplitude modulation (see fig. 2.6):

```
>> fm1=fmlin(256,0,0.5);  
>> am1=amgauss(256);  
>> sig1=am1.*fm1; plot(real(sig1));
```




Figure 2.6: Mono-component non-stationary signal with a linear frequency modulation and a gaussian amplitude modulation

By default, the signal is centered on the middle ($256/2=128$), and its spread is $T = 32$. If you want to center it at an other position t_0 , just replace `am1` by `amgauss(256,t0)`. A second example can be to multiply a pure frequency (constant frequency modulation) by a one-sided exponential window starting at $t=100$ (see fig. 2.7):

```
>> fm2=fmconst(256,0.2);  
>> am2=amexpo1s(256,100);  
>> sig2=am2.*fm2; plot(real(sig2));
```

As a third example of mono-component non-stationary signal, we can consider the M-file `doppler.m`: this function generates a modelization of the signal received by a fixed observer from a moving target emitting a pure frequency (see fig. 2.8).




Figure 2.7: Mono-component non-stationary signal with a constant frequency modulation and a one-sided exponential amplitude modulation

```
>> [fm3,am3]=doppler(256,200,4000/60,10,50);
>> sig3=am3.*fm3; plot(real(sig3));
```




Figure 2.8: Doppler signal

This example corresponds to a target (a car for instance) moving straightly at the speed of 50 m/s, and passing at 10 m from the observer (the radar!). The rotating frequency of the engine is 4000 revolutions per minute, and the sampling frequency of the radar is 200 Hz.

In order to have a more realistic modelization of physical signals, we may need to add some complex noise on these signals. To do so, two M-

files (`noisecg` and `noisecu`) of the Time-Frequency Toolbox are proposed : `noisecg.m` generates a complex white or colored gaussian noise, and `noisecu.m`, a complex white uniform noise. For example, if we add complex colored gaussian noise on the signal `sig1` with a signal to noise ratio of -10 dB (see fig. 2.9)

```
>> noise=noisecg(256,.8);  
>> sign=sigmerge(sig1,noise,-10); plot(real(sign));
```




Figure 2.9: Gaussian transient signal (`sig1`) embedded in a -10 dB colored gaussian noise

the deterministic signal `sig1` is now almost imperceptible from the noise.

2.7 What about multi-component non-stationary signals ?

The notion of instantaneous frequency implicitly assumes that, at each time instant, there exists only a single frequency component. A dual restriction applies to the group delay : the implicit assumption is that a given frequency is concentrated around a single time instant. Thus, if these assumptions are no longer valid, which is the case for most of the multi-component signals, the result obtained using the instantaneous frequency or the group delay is meaningless.

Example

For example, let us consider the superposition of two linear frequency modulations :

```
>> N=128; x1=fmlin(N,0,0.2); x2=fmlin(N,0.3,0.5);
>> x=x1+x2;
```

At each time instant t , an ideal time-frequency representation should represent two different frequencies with the same amplitude. The results obtained using the instantaneous frequency and the group delay are of course completely different, and therefore irrelevant (see fig. 2.10):

```
>> ifr=instfreq(x); subplot(211); plot(ifr);
>> fn=0:0.01:0.5; gd=sgrpdlay(x,fn);
>> subplot(212); plot(gd,fn);
```




Figure 2.10: Estimation of the instantaneous frequency (first plot) and group-delay (second plot) of a multi-component signal

So these one-dimensional representations, instantaneous frequency and group delay, are not sufficient to represent all the non-stationary signals. A further step has to be made towards two-dimensional mixed representations, jointly in time and in frequency. Even if no gain of information can be expected since it is all contained in the time or in the frequency representation, we can obtain a better structuring of this information, and an improvement in the intelligibility of the representation.

To have an idea of what can be made with a time-frequency decomposition, let us anticipate the following and have a look at the result obtained on this signal with the Short Time Fourier Transform (see fig. 2.11):

```
>> tfirstft(x);
```

Here two “time-frequency components” can be clearly seen, located around the locus of the two frequency modulations.




Figure 2.11: Squared modulus of the short-time Fourier transform of the previous multi-component non-stationary signal

Chapter 3

First class of solutions : the atomic decompositions

As we have seen in the previous chapter, the Fourier transform is not adapted to the analysis of non-stationary signals since it projects the signal on infinite waves (sinusoids) which are completely delocalized in time. The concepts of instantaneous frequency and group delay are also inherently unadapted to a large number of non-stationary signals, those containing more than one elementary component, and in particular noisy signals. Thus mono-dimensional solutions seem not to be sufficient, and one has to consider bi-dimensional functions (functions of the variables time and frequency).

A first class of such time-frequency representations is given by the *atomic decompositions* (also known as the *linear time-frequency representations*). To introduce this concept, we begin with the *short-time Fourier transform* which has a very intuitive interpretation.

3.1 The Short-Time Fourier Transform

3.1.1 Definition

In order to introduce time-dependency in the Fourier transform, a simple and intuitive solution consists in pre-windowing the signal $x(u)$ around a particular time t , calculating its Fourier transform, and doing that for each time instant t . The resulting transform, called the *short-time Fourier transform* (STFT, or *short-time spectrum*), is

$$F_x(t, \nu; h) = \int_{-\infty}^{+\infty} x(u) h^*(u - t) e^{-j2\pi\nu u} du$$

where $h(t)$ is a *short time analysis window* (see fig. 3.1) localized around $t = 0$ and $\nu = 0$. Because multiplication by the relatively short window




Figure 3.1: non-stationary signal $x(u)$ and the short-time window $h^*(u - t)$ centered at time t

$h^*(u - t)$ effectively suppresses the signal outside a neighborhood around the analysis time point $u = t$, the STFT is a "local" spectrum of the signal $x(u)$ around t . Provided that the short-time window is of finite energy, the STFT is invertible according to

$$x(t) = \frac{1}{E_h} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} F_x(u, \xi; h) h(t - u) e^{j2\pi t\xi} du d\xi,$$

with $E_h = \int_{-\infty}^{+\infty} |h(t)|^2 dt$. This relation expresses that the total signal can be decomposed as a weighted sum of elementary waveforms

$$h_{t,\nu}(u) = h(u - t) \exp[j2\pi\nu u]$$

which can be interpreted as "building blocks" or "*atoms*". Each atom is obtained from the window $h(t)$ by a translation in time and a translation in frequency (modulation). The corresponding transformation group of translations in both time and frequency is called the *Weyl-Heisenberg group*. Fig. 3.2 shows two such atoms corresponding to a gaussian window. The STFT may also be expressed in terms of signal and window spectra :

$$F_x(t, \nu; h) = \int_{-\infty}^{+\infty} X(\xi) H^*(\xi - \nu) \exp[j2\pi(\xi - \nu)t] d\xi$$

where X and H are respectively the Fourier transforms of x and h . Thus, the STFT $F_x(t, \nu; h)$ can be considered as the result of passing the signal




Figure 3.2: Time-frequency atoms: two atoms corresponding to a gaussian window. The STFT is a projection of the analyzed signal on such atoms which are relatively well localized in time and in frequency

$x(u)$ through a band-pass filter whose frequency response is $H^*(\xi - \nu)$, and is therefore deduced from a mother filter $H(\xi)$ by a translation of ν . So the STFT is similar to a bank of band-pass filters with constant bandwidth.

3.1.2 An example

Let us have a look at the result obtained by applying the STFT on a speech signal. The signal we consider is a speech signal containing the word 'GABOR', recorded on 338 points with a sampling frequency of 1 kHz (with respect to the Shannon criterion) (see fig. 3.3).

```
>> load gabor
>> time=0:337; subplot(211); plot(time,gabor);
>> dsp=fftshift(abs(fft(gabor)).^2);
>> freq=(-169:168)/338*1000; subplot(212); plot(freq,dsp);
```

We can not say from this representation what part of the word is responsible for that peak around 140 Hz.

Now if we look at the squared modulus of the STFT of this signal, using a hamming analysis window of 85 points, we can see some interesting features (the time-frequency matrix is loaded from the MAT-file because it takes a long time to be calculated ; we represent only the frequency domain where the signal is present) (see fig. 3.4):

```
>> contour(time,(0:127)/256*1000,tfr); grid;
```




Figure 3.3: Speech signal corresponding to the word 'GABOR'. Time signal (first plot) and its energy spectral density (second plot)

```
>> xlabel('Time [ms]'); ylabel('Frequency [Hz]');
>> title('Squared modulus of the STFT of the word GABOR');
```

The first pattern in the time-frequency plane, located between 30 ms and 60 ms, and centered around 150 Hz, corresponds to the first syllable 'GA'. The second pattern, located between 150 ms and 250 ms, corresponds to the last syllable 'BOR', and we can see that its mean frequency is decreasing from 140 Hz to 110 Hz with time. Harmonics corresponding to these two fundamental signals are also present at higher frequencies, but with a lower amplitude.

3.1.3 Some properties

- The STFT preserves frequency shifts and time shifts up to a modulation:

$$\begin{aligned}
 y(t) = x(t) e^{j2\pi\nu_0 t} &\Rightarrow F_y(t, \nu; h) = F_x(t, \nu - \nu_0; h) \\
 y(t) = x(t - t_0) &\Rightarrow F_y(t, \nu; h) = F_x(t - t_0, \nu; h) e^{j2\pi t_0 \nu}
 \end{aligned}$$

- Generalizing what has been said previously, the signal $x(t)$ can be reconstructed from its STFT with a synthesis window $g(t)$ different from the analysis window $h(t)$:

$$x(t) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} F_x(u, \xi; h) g(t - u) e^{j2\pi t \xi} du d\xi$$




Figure 3.4: Speech signal analyzed in the time-frequency plane

providing that the windows g and h validate the constraint

$$\int_{-\infty}^{+\infty} g(t) h^*(t) dt = 1.$$

3.1.4 Time-frequency resolution

The time resolution of the STFT can be obtained by considering for x a Dirac impulse :

$$x(t) = \delta(t - t_0) \Rightarrow F_x(t, \nu; h) = \exp[-j2\pi t_0 \nu] h(t - t_0).$$

Thus, the time resolution of the STFT is proportional to the effective duration of the analysis window h . Similarly, to obtain the frequency-resolution, we have to consider a complex sinusoid (a Dirac impulse in the frequency domain) :

$$x(t) = \exp[j2\pi \nu_0 t] \Rightarrow F_x(t, \nu; h) = \exp[-j2\pi t \nu_0] H(\nu - \nu_0).$$

So the frequency-resolution of the STFT is proportional to the effective bandwidth of the analysis window h . Consequently, for the STFT, we have a *trade-off* between time and frequency resolutions : on one hand, a good time resolution requires a short window $h(t)$; on the other hand, a good frequency

resolution requires a narrow-band filter i.e. a long window $h(t)$. But unfortunately, these wishes can not be simultaneously granted. This limitation is a consequence of the Heisenberg-Gabor inequality. Two instructive cases can be considered :

1. The first one corresponds to a perfect time resolution : the window $h(t)$ is chosen as a Dirac impulse :

$$h(t) = \delta(t) \Rightarrow F_x(t, \nu; h) = x(t) \exp[-j2\pi\nu t]$$

the STFT is perfectly localized in time, but does not provide any frequency resolution.

* *Example*: This can be computed easily using the Time-Frequency Toolbox : we consider for x a linear frequency modulation with a gaussian amplitude modulation (see fig. 3.5).

```
>> x=real(amgauss(128).*fmlin(128));
>> h=1;
>> tfrstft(x,1:128,128,h);
```




Figure 3.5: Perfect time resolution with the STFT, but with no frequency resolution : the window h is chosen as a Dirac impulse

The signal is perfectly localized in time (a section for a given frequency of the modulus of the STFT corresponds exactly to the modulus of the signal), but the frequency resolution is null.

2. The second is that of perfect frequency resolution, obtained with a constant window

$$h(t) = 1 \quad (H(\nu) = \delta(\nu)) \quad \Rightarrow \quad F_x(t, \nu; h) = X(\nu)$$

here the STFT reduces to the Fourier transform of $x(t)$, and does not provide any time resolution (see fig. 3.6).

```
>> h=ones(127,1);
>> tfrstft(x,1:128,128,h);
```




Figure 3.6: Perfect frequency resolution with the STFT : the window h is chosen as a constant

The result obtained for $F_x(t, \nu; h)$ is not exactly $X(\nu)$, because the window h has not an infinite duration. Thus, some side effects appear.

To illustrate the influence of the shape and length of the analysis window h , we consider two transient signals having the same gaussian amplitude and constant frequency, with different arrival times (using the M-file `atoms.m`):

```
>> sig=atoms(128,[45,.25,32,1;85,.25,32,1]);
```

Here is the result obtained with a Hamming analysis window of 65 points (see fig. 3.7) :

```
>> h=window(65,'hamming');
>> tfrstft(sig,1:128,128,h);
```




Figure 3.7: Two gaussian atoms analyzed by the STFT using a Hamming window h of 65 points : it is difficult to discriminate the two components in time

The frequency resolution is very good, but it is almost impossible to discriminate the two components in time. If we now consider a short Hamming window of 17 points (see fig. 3.8)

```
>> h=window(17,'hamming');
>> tfrstft(sig,1:128,128,h);
```

the frequency resolution is poorer, but the time resolution is sufficiently good to distinguish the two components. For more information on the choice of the window, see [Har78].

3.2 Time-scale analysis and the wavelet transform

Since the Wavelet Toolbox is fully dedicated to this problem, we will merely give here some basic definitions which are essential in the next part




Figure 3.8: Same gaussian atoms analyzed by the STFT using a Hamming window h of 17 points : frequency resolution is poorer, but the two components can be easily distinguished

to introduce the affine quadratic time-frequency distributions.

3.2.1 Definitions and interpretation

The idea of the *continuous wavelet transform* (CWT) is to project a signal x on a family of zero-mean functions (the *wavelets*) deduced from an elementary function (the *mother wavelet*) by translations and dilations:

$$T_x(t, a; \Psi) = \int_{-\infty}^{+\infty} x(s) \Psi_{t,a}^*(s) ds : \textit{Continuous Wavelet Transform}$$

where $\Psi_{t,a}(s) = |a|^{-1/2} \Psi\left(\frac{s-t}{a}\right)$. The variable a corresponds now to a *scale* factor, in the sense that taking $|a| > 1$ dilates the wavelet Ψ and taking $|a| < 1$ compresses Ψ . By definition, the wavelet transform is more a time-scale than a time-frequency representation. However, for wavelets which are well localized around a non-zero frequency ν_0 at scale $a = 1$, a time-frequency interpretation is possible thanks to the formal identification $\nu = \frac{\nu_0}{a}$.

The basic difference between the wavelet transform and the short-time Fourier transform is as follows: when the scale factor a is changed, the duration and the bandwidth of the wavelet are both changed but its shape




Figure 3.9: Time-scale atoms. The CWT is a projection of the analyzed signal on such atoms whose time duration is inversely proportional to the central frequency

remains the same. And in contrast to the STFT, which uses a single analysis window, the CWT uses short windows at high frequencies and long windows at low frequencies. This partially overcomes the resolution limitation of the STFT: the bandwidth B is proportional to ν , or

$$\frac{B}{\nu} = Q : \text{constant.}$$

We call it a *constant- Q analysis*. The CWT can also be seen as a filter bank analysis composed of band-pass filters with constant relative bandwidth.

3.2.2 Properties

- The wavelet transform is covariant by translation in time and scaling, which means that

$$y(t) = \sqrt{|a_0|} x(a_0(t - t_0)) \Rightarrow T_y(t, a; \Psi) = T_x(a_0^*(t - t_0), a/a_0; \Psi).$$

The corresponding group of transforms is called the *affine group* (to be compared to the Weyl-Heisenberg group).

- The signal x can be recovered from its continuous wavelet transform according to the formula

$$x(t) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} T_x(s, a; \Phi) \Psi_{s,a}(t) ds \frac{da}{a^2}$$

where Φ is the *synthesis wavelet*, if the following admissibility condition is verified by Φ and Ψ :

$$\int_{-\infty}^{+\infty} \Psi(\nu) \Phi^*(\nu) \frac{d\nu}{|\nu|} = 1.$$

- Time and frequency resolutions, like in the STFT case, are related via the Heisenberg-Gabor inequality. However, in the present case, these two resolutions depend on the frequency: the frequency resolution (resp. time resolution) becomes poorer (resp. better) as the analysis frequency grows.

3.3 Sampling considerations

3.3.1 The discrete STFT

To reduce the redundancy of the continuous STFT, we can sample it in the time-frequency plane. Since the atoms used can be deduced from the window $h(t)$ by translation in time and in frequency, it is natural to sample the STFT on a rectangular grid :

$$F_x[n, m; h] = F_x(nt_0, m\nu_0; h) = \int_{-\infty}^{+\infty} x(u) h^*(u - nt_0) \exp[-j2\pi m\nu_0 u] du$$

$m, n \in \mathcal{Z}$. The problem is then to choose the values of t_0 and ν_0 so as to minimize the redundancy without losing any information. For that, we must have

$$t_0 \times \nu_0 \leq 1.$$

Then, the atoms $h_{nt_0, m\nu_0}$ constitute a discrete over-sampled family of non orthonormal elements, which is called a *frame* : when $t_0 \times \nu_0 > 1$, the time-frequency plane is not sufficiently "covered" by the atoms $h_{nt_0, m\nu_0}$, i.e. there are "gaps" between adjacent atoms.

When $t_0 \times \nu_0 = 1$, the family of atoms $h_{nt_0, m\nu_0}$ can constitute an *orthonormal basis* for an appropriate choice of the window. But it can be shown that it is impossible to obtain such a basis with a window h which is well localized in time and in frequency (this property is known as the *Balian-Low obstruction* [Dau92]). Therefore, for a well localized window h (for example a gaussian window), the reconstruction formula will not be numerically stable.

In the discrete case, the reconstruction (synthesis) formula of the signal from the STFT is then given by

$$x(t) = \sum_n \sum_m F_x[n, m; h] g_{n,m}(t)$$

where $g_{n,m}(t) = g(t - nt_0) \exp[j2\pi m\nu_0 t]$.

This relation is valid provided that the sampling periods t_0 and ν_0 , the analysis window h and the synthesis window g are chosen such that

$$\frac{1}{\nu_0} \sum_n g\left(t + \frac{k}{\nu_0} - nt_0\right) h^*(t - nt_0) = \delta_k \quad \forall t$$

with δ_k defined as $\delta_0 = 1$ and $\delta_k = 0$ for $k \neq 0$. This condition is far more restrictive than the condition $\int_{-\infty}^{+\infty} g(t) h^*(t) dt = 1$ required in the continuous case.

For a sampled signal $x[n]$ whose sampling period is noted T , t_0 has to be chosen so that $t_0 = kT$, $k \in \mathcal{N}^*$. We then have the following analysis and synthesis formulae

$$F_x[n, m; h] = \sum_k x[k] h^*[k - n] \exp[-j2\pi mk] \quad \text{for } -\frac{1}{2} \leq m \leq \frac{1}{2} \quad (3.1)$$

$$x[k] = \sum_n \sum_m F_x[n, m; h] g[k - n] \exp[j2\pi mk]. \quad (3.2)$$

These two relations can be implemented efficiently using a Fast Fourier Transform (FFT) algorithm.

3.3.2 The Gabor Representation

The reconstruction (synthesis) formula of the STFT is given in the discrete case by

$$x(t) = \sum_n \sum_m F_x[n, m; h] g_{n,m}(t)$$

where $g_{n,m}(t) = g(t - nt_0) \exp[j2\pi m\nu_0 t]$ defines the *Gabor representation*. Originally, the synthesis window $g(t)$ was chosen by Gabor as a gaussian window, because it maximizes the concentration in the time-frequency plane. But now we speak of Gabor representation for any normalized window g .

The atoms $g_{n,m}(t)$ are called the *Gabor logons*, and the coefficients $F_x[n, m; h]$, noted $G_x[n, m]$ in the following, the *Gabor coefficients*. Each coefficient contains an information relative to the time-frequency content of the signal around the time-frequency location $(nt_0, m\nu_0)$. The logon $g_{n,m}$ is associated in the time-frequency plane to a rectangular unit area centered on $(nt_0, m\nu_0)$.

What about completeness of the Gabor logons $g_{n,m}(t)$? As we have seen before, a necessary but not sufficient condition is that $t_0 \nu_0 \leq 1$. At the critical sampling case $t_0 \nu_0 = 1$, the logons are linearly independent, but are not orthogonal in general (Balian-Low obstruction). This means that the Gabor coefficients $G_x[n, m]$ are not simply the projections of $x(t)$ onto the

corresponding logons $g_{n,m}(t)$ (i.e. the analysis and synthesis windows h and g can not be the same). A theoretical solution to this problem is obtained if the windows g and h are chosen biorthonormal, i.e. if they validate the *biorthonormal condition*

$$\int_{-\infty}^{+\infty} g_{n,m}(t) h_{n',m'}^*(t) dt = \delta_{n-n'} \delta_{m-m'}$$

Then the analysis formula given before (expression (3.1)) allows the computation of the Gabor coefficients, and the synthesis formula (expression (3.2)) the reconstruction of the signal $x(t)$ (to compute the *biorthonormal window* h associated to a given synthesis window g , one can use the *Zak transform* [AGT91] : this is the approach followed in the file `tfrgabor`, and the file `zak.m` computes this transform). From an implementation point of view, this solution is not fully satisfactory since the computation of the biorthonormal window h is numerically unstable. So in general, some degree of *oversampling* is considered ($t_0 \times \nu_0 < 1$), which introduces redundancy in the coefficients, in order to "smooth" the biorthonormal window h , for the sake of numerical stability. These considerations are closely connected to the theory of frames.

Example

Let us consider the Gabor coefficients of a linear chirp of $N1=256$ points at the critical sampling case, and for a gaussian window of $Ng=33$ points :

```
>> N1=256; Ng=33; Q=1; % degree of oversampling.
>> sig=fmlin(N1); g=window(Ng,'gauss'); g=g/norm(g);
>> [tfr,dgr,h]=tfrgabor(sig,16,Q,g);
```

(`tfrgabor` generates as first output the squared modulus of the Gabor representation, as second output the complex Gabor representation, and as third output the biorthonormal window). When we look at the biorthonormal window h (see fig. 3.10),

```
>> plot(h);
```

we can see how "bristling" this function is. The corresponding Gabor decomposition contains all the information about `sig`, but is not easy to interpret (see fig. 3.11) :

```
>> t=1:256; f=linspace(0,0.5,128); imagesc(t,f,tfr(1:128,:));
>> xlabel('Time'); ylabel('Normalized frequency'); axis('xy');
>> title('Squared modulus of the Gabor coefficients');
```




Figure 3.10: Biorthonormal window corresponding to the critical sampling case and to a gaussian synthesis window : numerically unsteady

If we now consider a degree of oversampling of $Q=4$ (there are four times more Gabor coefficients than signal samples), the biorthogonal function is then smoother (the greater Q , the closer h from g) (see fig. 3.12),

```
>> Q=4; [tfr,dgr,h]=tfrgabor(sig,32,Q,g);
>> plot(h);
```

and the Gabor representation is much more readable (see fig. 3.13):

```
>> imagesc(t,f,tfr(1:128,:));
>> xlabel('Time'); ylabel('Normalized frequency'); axis('xy');
>> title('Squared modulus of the Gabor coefficients');
```

3.3.3 The discrete wavelet transform

In the case of the wavelet transform, the natural way to sample the time-frequency plane is to take samples on the non-uniform grid (lattice) defined by

$$(t, a) = (nt_0 a_0^{-m}, a_0^{-m}); t_0 > 0, a_0 > 0; m, n \in \mathcal{Z}.$$

Then, the *discrete wavelet transform* (DWT) is defined as

$$T_x[n, m; \Psi] = a_0^{m/2} \int_{-\infty}^{+\infty} x(u) \Psi_{n,m}^*(u) du; m, n \in \mathcal{Z}$$

where $\Psi_{n,m}(u) = \Psi(a_0^m u - nt_0)$. The common choice ($a_0 = 2, t_0 = 1$) corresponds to a *dyadic sampling* of the time-frequency plane (one set of coefficients per octave) (see fig. 3.14). Thanks to such a sampling, it is now




Figure 3.11: Gabor representation of a chirp, at the critical sampling rate : we have as many coefficients in the time-frequency plane as in the signal (no redundancy)

possible to obtain for the family $\{\Psi_{n,m}(u) ; m, n \in \mathcal{Z}\}$ an *orthonormal basis* with a wavelet Ψ well localized in time and in frequency (the Balian-Low obstruction is no longer valid). This is strongly related to the multiresolution analysis theory (we will not develop it here ; see for more details the tutorial of the Wavelet Toolbox).

The main drawback of such a sampling is the loss of time-covariance. Indeed, a signal analyzed by the DWT will not have the same pattern on the dyadic grid whatever its initial position is.

As for the Gabor representation, a solution halfway between the over-complete family of wavelets $\Psi_{n,m}(u)$ used by the CWT and an orthonormal basis of wavelets obtained on the dyadic grid and for a particular choice of Ψ is given by the theory of frames (see [Dau92] for an overview of this theory with application to the wavelet transform).




Figure 3.12: Biorthonormal window h corresponding to an oversampling of $Q = 4$, and to a gaussian synthesis window g : the greater Q , the closer h from g

3.4 From atomic decompositions to energy distributions

Up to this point, we presented time-frequency representations that decompose the signal into elementary components, the atoms, well localized in time and in frequency. These representations were linear transforms of the signal.

Another approach to this problem, which will be developed in the next chapter, consists in distributing the energy of the signal along the two variables time and frequency. This gives rise to energy time-frequency distributions, which are naturally quadratic transforms of the signal.

We present in this section a natural transition between these two classes of solutions through the spectrogram (for the Weyl-Heisenberg group) and the scalogram (for the affine group).

3.4.1 The spectrogram

If we consider the squared modulus of the STFT, we obtain a spectral energy density of the locally windowed signal $x(u) h^*(u - t)$:

$$S_x(t, \nu) = \left| \int_{-\infty}^{+\infty} x(u) h^*(u - t) e^{-j2\pi\nu u} du \right|^2.$$

This defines the *spectrogram*, which is a real-valued and non-negative distribution. Since the window h of the STFT is assumed of unit energy, the




Figure 3.13: Gabor representation of the same chirp, but with a degree of oversampling of 4 : some redundancy improve the readability of the representation

spectrogram satisfies the global energy distribution property :

$$\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} S_x(t, \nu) dt d\nu = E_x.$$

Thus, we can interpret the spectrogram as a measure of the energy of the signal contained in the time-frequency domain centered on the point (t, ν) and whose shape is independent of this localization.

- *Properties*

- Time and frequency covariance

A direct consequence of the definition of the spectrogram is that it preserves time and frequency shifts :

$$\begin{aligned} y(t) = x(t - t_0) &\Rightarrow S_y(t, \nu) = S_x(t - t_0, \nu) \\ y(t) = x(t) \exp [j2\pi\nu_0 t] &\Rightarrow S_y(t, \nu) = S_x(t, \nu - \nu_0). \end{aligned}$$

Thus, the spectrogram is an element of the class of quadratic time-frequency distributions that are covariant by translation in time and in frequency. This class, developed in the next chapter, is called the *Cohen's class*.




Figure 3.14: Sampling of the time-frequency plane. Different forms of sampling: Shannon, Fourier, Gabor, Wavelet

– Time-frequency resolution

The spectrogram being the squared magnitude of the STFT, it is obvious that the time-frequency resolution of the spectrogram is limited exactly as it is for the STFT. In particular, it exists again a trade-off between time resolution and frequency resolution. This poor resolution property is the main drawback of this representation.

– Interference structure

As it is a quadratic (or bilinear) representation, the spectrogram of the sum of two signals is not the sum of the two spectrograms (*quadratic superposition principle*):

$$y(t) = x_1(t) + x_2(t) \Rightarrow S_y(t, \nu) = S_{x_1}(t, \nu) + S_{x_2}(t, \nu) + 2\Re\{S_{x_1, x_2}(t, \nu)\}$$

where $S_{x_1,x_2}(t,\nu) = F_{x_1}(t,\nu)F_{x_2}^*(t,\nu)$ is the cross-spectrogram and \Re denotes the real part. Thus, as every quadratic distribution, the spectrogram presents interference terms, given by $S_{x_1,x_2}(t,\nu)$. However, one can show [Hla91] that these interference terms are restricted to those regions of the time-frequency plane where the auto-spectrograms $S_{x_1}(t,\nu)$ and $S_{x_2}(t,\nu)$ overlap. Thus, if the signal components $x_1(t)$ and $x_2(t)$ are sufficiently distant so that their spectrograms do not overlap significantly, then the interference term will nearly be identically zero. This property, which is a practical advantage of the spectrogram, is in fact a consequence of the spectrogram's poor resolution.

- *Examples*

To illustrate the resolution trade-off of the spectrogram and its interference structure, we consider a two-component signal composed of two parallel chirps, and we analyze it with the M-file `tfrsp.m` of the Time-Frequency Toolbox (the M-file `specgram.m` of the Signal Processing Toolbox is equivalent, except that `tfrsp.m` offers the possibility to change the analysis window) (see fig. 3.15 and fig. 3.16).

```
>> sig=fmlin(128,0,0.4)+fmlin(128,0.1,0.5);
>> h1=window(23,'gauss');
>> figure(1); tfrsp(sig,1:128,128,h1);
>> h2=window(63,'gauss');
>> figure(2); tfrsp(sig,1:128,128,h2);
```

In these two cases, the two FM components of the signal are not sufficiently distant to have distinct spectrograms, whatever the window length is. Consequently, interference terms are present, and disturb the readability of the time-frequency representation. If we consider more distant components (see fig. 3.17 and fig. 3.18),

```
>> sig=fmlin(128,0,0.3)+fmlin(128,0.2,0.5);
>> h1=window(23,'gauss');
>> figure(1); tfrsp(sig,1:128,128,h1);
>> h2=window(63,'gauss');
>> figure(2); tfrsp(sig,1:128,128,h2);
```

the two auto-spectrograms do not overlap and no interference term appear. We can also see the effect of a short window (`h1`) and a long window (`h2`) on the time-frequency resolution. In the present case, the




Figure 3.15: Spectrogram of two parallel chirps, using a short gaussian analysis window : cross-terms are present between the two FM components

long window h_2 is preferable since as the frequency progression is not very fast, the quasi-stationary assumption will be correct over h_2 (so time resolution is not as important as frequency resolution in this case) and the frequency resolution will be quite good ; whereas if the window is short (h_1), the time resolution will be good, which is not very useful, and the frequency resolution will be poor.

3.4.2 The scalogram

A similar distribution to the spectrogram can be defined in the wavelet case. Since the continuous wavelet transform behaves like an orthonormal basis decomposition, it can be shown that it preserves energy :

$$\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} |T_x(t, a; \Psi)|^2 dt \frac{da}{a^2} = E_x$$

where E_x is the energy of x . This leads us to define the *scalogram* of x as the squared modulus of the continuous wavelet transform. It is an energy distribution of the signal in the time-scale plane, associated with the measure $dt \frac{da}{a^2}$.

As for the wavelet transform, time and frequency resolutions of the scalogram are related via the Heisenberg-Gabor principle: time and frequency




Figure 3.16: Spectrogram of two parallel chirps, using a long gaussian analysis window : cross-terms are still present, due to the too small distance in the time-frequency plan between the FM components

resolutions depend on the considered frequency. To illustrate this point, we represent the scalograms of two different signals. The M-file `tfrscalogram` generates this representation. The chosen wavelet is a Morlet wavelet of 12 points. The first signal is a Dirac pulse at time $t_0 = 64$:

```
>> sig1=anapulse(128);
>> tfrscalogram(sig1,1:128,6,0.05,0.45,128,1);
```

Figure 3.19 shows that the influence of the behavior of the signal around $t = t_0$ is limited to a cone in the time-scale plane: it is "very" localized around t_0 for small scales (large frequencies), and less and less localized as the scale increases (as the frequency decreases).

The second signal is the sum of two sinusoids of different frequencies (see fig. 3.20) :

```
>> sig2=fmconst(128,.15)+fmconst(128,.35);
>> tfrscalogram(sig2,1:128,6,0.05,0.45,128,1);
```

Here again, we notice that the frequency resolution is clearly a function of the frequency: it increases with ν .




Figure 3.17: Spectrogram of two more distant parallel chirps, using a short gaussian analysis window

The interference terms of the scalogram, as for the spectrogram, are also restricted to those regions of the time-frequency plane where the corresponding auto-scalograms (signal terms) overlap. Hence, if two signal components are sufficiently far apart in the time-frequency plane, their cross-scalogram will be essentially zero.

3.4.3 Conclusion

Through this chapter, we presented a first class of time-frequency distributions of non-stationary signals. These distributions decompose the signal on a basis of elementary signals (the atoms) which have to be well localized in time and in frequency. Two well known examples of such decompositions are the short-time Fourier transform and the wavelet transform. After having considered their properties, we discussed their formulation in the discrete case. Finally, we presented a natural transition from this class of representations to the class of energy distributions.




Figure 3.18: Spectrogram of two parallel chirps, using a long gaussian analysis window




Figure 3.19: Morlet scalogram of a Dirac impulse at time $t = 64$: time resolution depends on the considered frequency (or scale)




Figure 3.20: Morlet scalogram of two simultaneous complex sinusoids : frequency resolution depends on the considered frequency (or scale)

Chapter 4

Second class of solutions : the energy distributions

In contrast with the linear time-frequency representations which decompose the signal on elementary components (the atoms), the purpose of the energy distributions is to distribute the *energy* of the signal over the two description variables : time and frequency.

The starting point is that since the energy of a signal x can be deduced from the squared modulus of either the signal or its Fourier transform,

$$E_x = \int_{-\infty}^{+\infty} |x(t)|^2 dt = \int_{-\infty}^{+\infty} |X(\nu)|^2 d\nu, \quad (4.1)$$

we can interpret $|x(t)|^2$ and $|X(\nu)|^2$ as energy densities, respectively in time and in frequency. It is then natural to look for a *joint time and frequency* energy density $\rho_x(t, \nu)$, such that

$$E_x = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \rho_x(t, \nu) dt d\nu, \quad (4.2)$$

which is an intermediary situation between those described by (4.1). As the energy is a quadratic function of the signal, the time-frequency energy distributions will be in general quadratic representations.

Two other properties that an energy density should satisfy are the following *marginal properties* :

$$\int_{-\infty}^{+\infty} \rho_x(t, \nu) dt = |X(\nu)|^2 \quad (4.3)$$

$$\int_{-\infty}^{+\infty} \rho_x(t, \nu) d\nu = |x(t)|^2, \quad (4.4)$$

which mean that if we integrate the time-frequency energy density along one variable, we obtain the energy density corresponding to the other variable.

The main references for this chapter are [Fla93], [Coh89], [Aug91], [Hla91] and [HBB92].

4.1 The Cohen's class

Since there is much more than one distribution satisfying properties (4.2), (4.3) and (4.4), we can impose additional constraints on ρ_x so that this distribution satisfies other desirable properties. Among these, the covariance principles are of fundamental importance. The *Cohen's class*, to which is dedicated this section, and whose definition can be found in subsection 4.1.2, is the class of time-frequency energy distributions *covariant by translations in time and in frequency* [Coh89].

The spectrogram, that we considered in the previous part, is an element of the Cohen's class since it is quadratic, time- and frequency- covariant, and preserves energy (property (4.2)). However, taking the squared modulus of an atomic decomposition is only a restrictive possibility to define a quadratic representation, and this definition presents the drawback that the marginal properties (4.3) and (4.4) are not satisfied.

4.1.1 The Wigner-Ville distribution

Definition

A time-frequency energy distribution which is particularly interesting is the *Wigner-Ville distribution* (WVD) defined as :

$$W_x(t, \nu) = \int_{-\infty}^{+\infty} x(t + \tau/2) x^*(t - \tau/2) e^{-j2\pi\nu\tau} d\tau, \quad (4.5)$$

or equivalently as

$$W_x(t, \nu) = \int_{-\infty}^{+\infty} X(\nu + \xi/2) X^*(\nu - \xi/2) e^{j2\pi\xi t} d\xi.$$

This distribution satisfies a large number of desirable mathematical properties, as summarized in the next sub-section. In particular, the WVD is always real-valued, it preserves time and frequency shifts and satisfies the marginal properties.

An interpretation of this expression can be found in terms of probability density: expression (4.5) is the Fourier transform of an acceptable form of characteristic function for the distribution of the energy.

Before looking at the theoretical properties of the WVD, let us see what we obtain on two particular synthetic signals.

- *Example 1*: The first signal is the academic linear chirp signal that we already considered. The WVD is available thanks to the M-file `tfrwv.m` of the Time-Frequency Toolbox (see fig. 4.1).

```
>> sig=fmlin(256);
>> tfrwv(sig);
```




Figure 4.1: Wigner-Ville distribution of a linear chirp signal : almost perfect localization in the time-frequency plane

If we choose a 3-dimension plot to represent it, we can see that the WVD can take negative values, and that the localization obtained in the time-frequency plane for this signal is almost perfect.

- *Example 2*: When a car goes in front of an observer with a constant speed, the signal heard by this person from the engine changes with time: the main frequency decreases (at a first level of approximation) from one value to another. This phenomenon, known as the *doppler effect*, expresses the dependence of the frequency received by an observer from a transmitter on the relative speed between the observer and the transmitter. The corresponding signal can be generated thanks to the M-file `doppler.m` of the Time-Frequency Toolbox. Here is an example of such a signal (see fig. 4.2):

```

>> [fm,am,iflaw]=doppler(256,50,13,10,200);
>> sig=am.*fm;
>> tfrwv(sig);

```




Figure 4.2: WVD of a doppler signal : many interference terms are present, due to the bilinearity of the distribution

Looking at this time-frequency distribution, we notice that the energy is not distributed as we could expect for this signal. Although the signal term is well localized in the time-frequency plane, numerous other terms (the interference terms, due to the bilinearity of the WVD) are present at positions in time and frequency where the energy should be null. We will see earlier how to get rid of these terms.

Properties

Here is a list of the main properties of the WVD [Fla93].

1. *Energy conservation*: by integrating the WVD of x all over the time-frequency plane, we obtain the energy of x :

$$E_x = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} W_x(t, \nu) dt d\nu$$

2. *Marginal properties*: the energy spectral density and the instantaneous power can be obtained as marginal distributions of W_x :

$$\int_{-\infty}^{+\infty} W_x(t, \nu) dt = |X(\nu)|^2$$

$$\int_{-\infty}^{+\infty} W_x(t, \nu) d\nu = |x(t)|^2$$

3. *Real-valued*:

$$W_x(t, \nu) \in \mathcal{R}, \forall t, \nu$$

4. *Translation covariance*: the WVD is time and frequency covariant:

$$y(t) = x(t - t_0) \Rightarrow W_y(t, \nu) = W_x(t - t_0, \nu)$$

$$y(t) = x(t)e^{j2\pi\nu_0 t} \Rightarrow W_y(t, \nu) = W_x(t, \nu - \nu_0)$$

5. *Dilation covariance*: the WVD also preserves dilations:

$$y(t) = \sqrt{k} x(kt) ; k > 0 \Rightarrow W_y(t, \nu) = W_x(kt, \frac{\nu}{k})$$

6. *Compatibility with filterings*: it expresses the fact that if a signal y is the convolution of x and h (i.e. the output of filter h whose input is x), the WVD of y is the time-convolution between the WVD of h and the WVD of x :

$$y(t) = \int_{-\infty}^{+\infty} h(t-s) x(s) ds \Rightarrow W_y(t, \nu) = \int_{-\infty}^{+\infty} W_h(t-s, \nu) W_x(s, \nu) ds$$

7. *Compatibility with modulations*: this is the dual property of the previous one: if y is the modulation of x by a function m , the WVD of y is the frequency-convolution between the WVD of x and the WVD of m :

$$y(t) = m(t) x(t) \Rightarrow W_y(t, \nu) = \int_{-\infty}^{+\infty} W_m(t, \nu - \xi) W_x(t, \xi) d\xi$$

8. *Wide-sense support conservation*: if a signal has a compact support in time (respectively in frequency), then its WVD also has the same compact support in time (respectively in frequency):

$$x(t) = 0, |t| > T \Rightarrow W_x(t, \nu) = 0, |t| > T$$

$$X(\nu) = 0, |\nu| > B \Rightarrow W_x(t, \nu) = 0, |\nu| > B$$

9. *Unitarity*: the unitarity property expresses the conservation of the scalar product from the time-domain to the time-frequency domain (apart from the squared modulus):

$$\left| \int_{-\infty}^{+\infty} x(t) y^*(t) dt \right|^2 = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} W_x(t, \nu) W_y^*(t, \nu) dt d\nu.$$

This formula is also known as the Moyal's formula.

10. *Instantaneous frequency*: the instantaneous frequency of a signal x can be recovered from the WVD as its first order moment (or center of gravity) in frequency:

$$f_x(t) = \frac{\int_{-\infty}^{+\infty} \nu W_{x_a}(t, \nu) d\nu}{\int_{-\infty}^{+\infty} W_{x_a}(t, \nu) d\nu}$$

where x_a is the analytic signal associated to x .

11. *Group delay*: in a dual way, the group delay of x can be obtained as the first order moment in time of its WVD:

$$t_x(\nu) = \frac{\int_{-\infty}^{+\infty} t W_{x_a}(t, \nu) dt}{\int_{-\infty}^{+\infty} W_{x_a}(t, \nu) dt}$$

12. *Perfect localization on linear chirp signals*:

$$x(t) = e^{j2\pi\nu_x(t)t} \text{ with } \nu_x(t) = \nu_0 + 2\beta t \Rightarrow W_x(t, \nu) = \delta(\nu - (\nu_0 + \beta t)).$$

Interferences

As the WVD is a bilinear function of the signal x , the *quadratic superposition principle* applies:

$$W_{x+y}(t, \nu) = W_x(t, \nu) + W_y(t, \nu) + 2\Re\{W_{x,y}(t, \nu)\}$$

where

$$W_{x,y}(t, \nu) = \int_{-\infty}^{+\infty} x(t + \tau/2) y^*(t - \tau/2) e^{-j2\pi\nu\tau} d\tau$$

is the cross-WVD of x and y . This can be easily generalized to N components, but for the sake of clarity, we will only consider the two-component case.

Unlike the spectrogram interference terms, the WVD interference terms will be non-zero regardless of the time-frequency distance between the two signal terms. These interference terms are troublesome since they may overlap with auto-terms (signal terms) and thus make it difficult to visually

interpret the WVD image. However, it appears that these terms must be present or the good properties of the WVD (marginal properties, instantaneous frequency and group delay, localization, unitarity ...) cannot be satisfied. Actually, there is a trade-off between the quantity of interferences and the number of good properties.

o *Interference geometry*

The rule of interference construction of the WVD can be summarized as follows: two points of the time-frequency plane interfere to create a contribution on a third point which is located at their geometrical midpoint. Besides, these interference terms oscillate perpendicularly to the line joining the two points interfering, with a frequency proportional to the distance between these two points.

This can be seen on the following example: we consider two atoms in the time-frequency plane, analyzed by the WVD, whose relative distance is increasing from one realization to the other, and then decreasing. The WVDs were calculated and saved on the file `movwv2at.mat`. We load them and run the sequence using the function `movie` (see fig. 4.3):

```
>> load movwv2at
>> clf; movie(M,10);
```

We can notice, from this movie, the evolution of the interferences when the distance between the two interfering terms changes, and in particular the change in the direction of the oscillations.

Pseudo-WVD

The definition (4.5) requires the knowledge of the quantity

$$q_x(t, \tau) = x(t + \tau/2) x^*(t - \tau/2)$$

from $\tau = -\infty$ to $\tau = +\infty$, which can be a problem in practice. That is why we often replace $q_x(t, \tau)$ in (4.5) by a windowed version of it, leading to the new distribution:

$$PW_x(t, \nu) = \int_{-\infty}^{+\infty} h(\tau) x(t + \tau/2) x^*(t - \tau/2) e^{-j2\pi\nu\tau} d\tau$$

where $h(t)$ is a regular window. This distribution is called the *pseudo Wigner-Ville distribution* (noted pseudo-WVD or PWVD in the following). This windowing operation is equivalent to a frequency smoothing of the WVD since

$$PW_x(t, \nu) = \int_{-\infty}^{+\infty} H(\nu - \xi) W_x(t, \xi) d\xi$$




Figure 4.3: Structure of the interferences between 2 components with different locations in time and frequency : we can notice the change in the direction of the oscillations, as well as the change in the period of these oscillations

where $H(\nu)$ is the Fourier transform of $h(t)$. Thus, because of their oscillating nature, the interferences will be attenuated in the pseudo-WVD compared to the WVD. However, the consequence of this improved readability is that many properties of the WVD are lost : the marginal properties, the unitarity, and also the frequency-support conservation ; the frequency-widths of the auto-terms are increased by this operation.

* *Example*: The M-file `tfrpwv.m` calculates the pseudo-WVD of a signal, with the possibility to change the length and shape of the smoothing window. If we consider a signal composed of four gaussian atoms (obtained thanks to `atoms.m`), each localized at a corner of a rectangle,

```
>> sig=atoms(128,[32,.15,20,1;96,.15,20,1;...
                32,.35,20,1;96,.35,20,1]);
```

and compute its WVD (see fig. 4.4)

```
>> tfrwv(sig);
```




Figure 4.4: WVD of 4 gaussian atoms : many interferences are present

we can see the four signal terms, along with six interference terms (two of them are superimposed). If we now compute the pseudo-WVD (see fig. 4.5),

```
>> tfrpwv(sig);
```

we can note the important attenuation of the interferences oscillating perpendicularly to the frequency axis, and in return the spreading in frequency of the signal terms.

Sampling the WVD ; the analytic signal

Because of the quadratic nature of the WVD, its sampling has to be done with care. Let us write it as follows :

$$W_x(t, \nu) = 2 \int_{-\infty}^{+\infty} x(t + \tau) x^*(t - \tau) e^{-j4\pi\nu\tau} d\tau$$

If we sample x with a period T_e , write $x[n] = x(nT_e)$, and evaluate the WVD at the sampling points nT_e in time, we obtain a discrete-time continuous-frequency expression of it :

$$W_x[n, \nu) = 2 T_e \sum_k x[n + k] x^*[n - k] e^{-j4\pi\nu k}.$$

As this expression is periodic in frequency with period $\frac{1}{2T_e}$ (contrary to period $\frac{1}{T_e}$ obtained for the Fourier transform of a signal sampled at the Nyquist




Figure 4.5: The frequency-smoothing operated by the pseudo-WVD attenuates the interferences oscillating perpendicularly to the frequency axis

rate), the discrete version of the WVD may be affected by a spectral aliasing, in particular if the signal x is real-valued and sampled at the Nyquist rate. Two alternatives to this problem can be found. The first one consists in oversampling the signal by a factor of at least 2, and the second one in using the analytic signal. Indeed, as its bandwidth is half the one of the real signal, the aliasing will not take place in the useful spectral domain $[0, 1/2]$ of this signal. This second solution presents another advantage: since the spectral domain is divided by two, the number of components in the time-frequency plane is also divided by two. Consequently, the number of interference terms decreases significantly. To illustrate this phenomenon, we consider the WVD of the real part of a signal composed of two atoms (see fig. 4.6) :

```
>> sig=atoms(128,[32,0.15,20,1;96,0.32,20,1]);
>> tfrwv(real(sig));
```

We can see that four signal terms are present instead of two, due to the spectral aliasing. Besides, because of the components located at negative frequencies (between $-1/2$ and 0), additional interference terms are present. If we now consider the WVD of the same signal, but in its complex analytic form (see fig. 4.7),

```
>> tfrwv(sig);
```




Figure 4.6: WVD of a real signal composed of 2 gaussian atoms : when the analytic signal is not considered, spectral aliasing and additional interferences appear in the time-frequency plane

the aliasing effect has disappeared, as well as the terms corresponding to interferences between negative- and positive- frequency components.

4.1.2 The Cohen's class

Presentation

Among the desirable properties of an energy time-frequency distribution, two of them are of particular importance: *time and frequency covariance*. Indeed, these properties guaranty that, if the signal is delayed in time and modulated, its time-frequency distribution is translated of the same quantities in the time-frequency plane. It has been shown that the class of energy time-frequency distributions verifying these covariance properties possesses the following general expression :

$$C_x(t, \nu; f) = \int \int \int_{-\infty}^{+\infty} e^{j2\pi\xi(s-t)} f(\xi, \tau) x(s+\tau/2) x^*(s-\tau/2) e^{-j2\pi\nu\tau} d\xi ds d\tau,$$

where $f(\xi, \tau)$ is a two-dimensional function called the *parameterization function*. This class of distributions is known as the *Cohen's class*, which can also




Figure 4.7: WVD of the previous signal, but in its analytic form

be written :

$$C_x(t, \nu; \Pi) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Pi(s - t, \xi - \nu) W_x(s, \xi) ds d\xi, \quad (4.6)$$

where

$$\Pi(t, \nu) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(\xi, \tau) e^{-j2\pi(\nu\tau + \xi t)} dt d\nu$$

is the two-dimensional Fourier transform of the parameterization function f . This class is of significant importance since it includes a large number of the existing time-frequency energy distributions. Of course, the WVD is the element of the Cohen's class for which the function Π is a double Dirac : $\Pi(t, \nu) = \delta(t) \delta(\nu)$, i.e. $f(\xi, \tau) = 1$.

In the case where Π is a smoothing function, expression (4.6) allows one to interpret C_x as a smoothed version of the WVD ; consequently, such a distribution will attenuate in a particular way the interferences of the WVD.

Before considering different kinds of smoothing functions Π , let us point out the different advantages of such a unified formulation :

1. by specifying the parameterization function f arbitrarily, it is possible to obtain most of the known energy distributions ;
2. it is easy to convert a constraint that we wish for the distribution in an admissibility condition for the parameterization function ;

3. it is possible, by using such admissibility arguments, to check *a priori* the properties of a particular definition, or to construct a class of solutions according to a specified schedule of conditions.

Coupled smoothing

If we look at the Moyal's formula (property 9. see page 62), it is easy to express the spectrogram as a smoothing of the WVD :

$$S_x(t, \nu) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} W_h(s - t, \xi - \nu) W_x(s, \xi) ds d\xi. \quad (4.7)$$

Thus, the spectrogram is the element of the Cohen's class for which $\Pi(s, \xi)$ is the WVD of the window h . This new formulation provides us with another interpretation of the embarrassing trade-off between the time and frequency-resolutions of the spectrogram : if we choose a short window h , the smoothing function will be narrow in time and wide in frequency, leading to a good time resolution but bad frequency resolution ; and vice-versa.

Separable smoothing

The problem with the previous smoothing function $\Pi(s, \xi) = W_h(s, \xi)$ is that it is controlled only by the short-time window $h(t)$. If we add a degree of freedom by considering a separable smoothing function

$$\Pi(t, \nu) = g(t) H(-\nu)$$

(where $H(\nu)$ is the Fourier transform of a smoothing window $h(t)$), we allow a progressive and independent control, in both time and frequency, of the smoothing applied to the WVD. The obtained distribution

$$SPW_x(t, \nu) = \int_{-\infty}^{+\infty} h(\tau) \int_{-\infty}^{+\infty} g(s - t) x(s + \tau/2) x^*(s - \tau/2) ds e^{-j2\pi\nu\tau} d\tau$$

is known as the *smoothed-pseudo Wigner-Ville distribution* (noted smoothed-pseudo-WVD or SPWVD). The previous compromise of the spectrogram between time and frequency- resolutions is now replaced by a compromise between the joint time-frequency resolution and the level of the interference terms : the more you smooth in time and/or frequency, the poorer the resolution in time and/or frequency.

Note that if we only consider a smoothing in frequency i.e. if $g(t) = \delta(t)$, we obtain the pseudo-WVD.

* *Example*: The signal that we consider here is composed of two components: the first one is a complex sinusoid (normalized frequency 0.15) and the second one is a Gaussian signal shifted in time and frequency :

```
>> sig=fmconst(128,.15) + amgauss(128).*fmconst(128,0.4);
```

If we display the WVD, the pseudo-WV and the smoothed-pseudo-WVD of this signal (see fig. 4.8, fig. 4.9 and fig. 4.10),

```
>> tfrwv(sig);
>> tfrpwv(sig);
>> tfrspwv(sig);
```




Figure 4.8: WVD of a signal composed of a Gaussian atom and a complex sinusoid. Interferences are present between the two components

we can make the following remarks : from the WVD, we can see the two signal terms located at the right positions in the time-frequency plane, as well as the interference terms between them. As these interference terms oscillate globally perpendicularly to the time-axis, the frequency smoothing done by the pseudo-WVD degrades the frequency resolution without really attenuating the interferences. On the other hand, the time-smoothing carried out by the smoothed-pseudo-WVD considerably reduces these interferences ; and as the time resolution is not of fundamental importance here, this representation is suitable for this signal.




Figure 4.9: Pseudo-WVD of the same signal : the frequency smoothing done by the pseudo-WVD degrades the frequency resolution without really attenuating the interferences

An interesting property of the smoothed-pseudo WVD is that it allows a continuous passage from the spectrogram to the WVD, under the condition that the smoothing functions g and h are gaussian. The time-bandwidth product then goes from 1 (spectrogram) to 0 (WVD), with an independent control of the time and frequency resolutions. This is clearly illustrated by the function `movsp2wv.m`, which considers different transitions, on a signal composed of four atoms. To visualize these snapshots, load the mat-file `movsp2wv` (obtained by running `movsp2wv.m`; but as it takes a long time to run, we saved the result in a mat file) and run `movie` (see fig. 4.11):

```
>> load movsp2wv
>> clf; movie(M,10);
```

This movie shows the effect of a (time/frequency) smoothing on the interferences and on the resolutions: the WVD gives the best resolutions (in time and in frequency), but presents the most important interferences, whereas the spectrogram gives the worst resolutions, but with nearly no interferences; and the smoothed-pseudo WVD allows to choose the best compromise between these two extremes.




Figure 4.10: Smoothed-pseudo-WVD of the same signal : the time-smoothing carried out by the smoothed-pseudo-WVD considerably reduces these interferences

4.1.3 Link with the narrow-band ambiguity function

Definition and properties

A function of particular interest, especially in the field of radar signal processing, is the *narrow-band ambiguity function* (noted AF), defined as

$$A_x(\xi, \tau) = \int_{-\infty}^{+\infty} x(s + \tau/2) x^*(s - \tau/2) e^{-j2\pi\xi s} ds.$$

This function, also known as the (symmetric) *Sussman ambiguity function*, is a measure of the time-frequency correlation of a signal x , i.e. the degree of similarity between x and its translated versions in the time-frequency plane. Unlike the variables ' t ' and ' ν ' which are "absolute" time and frequency coordinates, the variables ' τ ' and ' ξ ' are "relative" coordinates (respectively called *delay* and *doppler*).

The AF is generally complex-valued, and satisfies the Hermitian even symmetry :

$$A_x(\xi, \tau) = A_x^*(-\xi, -\tau).$$

An important relation exists between the narrow-band ambiguity function and the WVD, which says that the ambiguity function is the two-dimensional




Figure 4.11: Different transitions from the spectrogram to the WVD, using the smoothed-pseudo-WVD. The signal is composed of 4 gaussian atoms

Fourier transform of the WVD :

$$A_x(\xi, \tau) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} W_x(t, \nu) e^{j2\pi(\nu\tau - \xi t)} dt d\nu.$$

Thus, the AF is the dual of the WVD in the sense of the Fourier transform. Consequently, for the AF, a dual property corresponds to nearly all the properties of the WVD. Among these properties, we will restrict ourselves to only three of them, which are important for the following :

- Marginal properties

The temporal and spectral auto-correlations are the cuts of the AF along the τ -axis and ξ -axis respectively :

$$r_x(\tau) = A_x(0, \tau) \text{ and } R_x(\xi) = A_x(\xi, 0).$$

The energy of x is the value of the AF at the origin of the (ξ, τ) -plane, which corresponds to its maximum value :

$$|A_x(\xi, \tau)| \leq A_x(0, 0) = E_x, \forall \xi, \tau.$$

- TF-shift invariance

Shifting a signal in the time-frequency plane leaves its AF invariant apart from a phase factor (modulation):

$$y(t) = x(t - t_0) e^{j2\pi\nu_0 t} \Rightarrow A_y(\xi, \tau) = A_x(\xi, \tau) e^{j2\pi(\nu_0\tau - t_0\xi)}$$

- Interference geometry

In the case of a multi-component signal, the elements of the AF corresponding to the signal components (denoted as the AF-signal terms) are mainly located around the origin, whereas the elements corresponding to interferences between the signal components (AF-interference terms) appear at a distance from the origin which is proportional to the time-frequency distance between the involved components. This can be noticed on a simple example:

* *Example:* The M-file `ambifunb.m` of the TF Toolbox implements the narrow-band ambiguity function. We apply it on a signal composed of two linear FM signals with gaussian amplitudes:

```
>> N=64; sig1=fmlin(N,0.2,0.5).*amgauss(N);
>> sig2=fmlin(N,0.3,0).*amgauss(N);
>> sig=[sig1;sig2];
```

Let us first have a look at the WVD (see fig. 4.12):

```
>> tfrwv(sig);
```

We have two distinct signal terms, and some interferences oscillating in the middle. If we look at the ambiguity function of this signal (see fig. 4.13),

```
>> ambifunb(sig);
```

we have around the origin (in the middle of the image) the AF-signal terms, whereas the AF-interference terms are located away from the origin. Thus, applying a 2-D low pass filtering around the origin on the ambiguity function, and returning to the WVD by 2-D Fourier transform will attenuate the interference terms. Actually, this 2-D filtering is operated, in the general expression of the Cohen's class, by the parameterization function f , as we discuss it now.




Figure 4.12: WVD of 2 chirps with gaussian amplitudes and different slopes

New interpretation of the Cohen's class

The dual expression of the Cohen's class formulation (expression (4.6)) in terms of AF writes

$$C_x(t, \nu; f) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} f(\xi, \tau) A_x(\xi, \tau) e^{-j2\pi(\nu\tau + \xi t)} d\xi d\tau \quad (4.8)$$

(recall that f is the two-dimensional Fourier transform of Π). This expression is very instructive about the role played by the parameterization function $f(\xi, \tau)$. Indeed, f acts as a weighting function that tries to let the signal terms unchanged, and to reject the interference terms. Actually, the change from the time-frequency plane to the ambiguity plane allows a precise characterization of the weighting function f , and thus of the smoothing function $\Pi(t, \nu)$.

For example, the WVD corresponds to a constant parameterization function: $f(\xi, \tau) = 1, \forall \xi, \tau$: no difference is made between the different regions of the ambiguity plane. For the spectrogram, $f(\xi, \tau) = A_h^*(\xi, \tau)$: the ambiguity function of the window h determines the shape of the weighting function. And for the smoothed-pseudo-WVD, we have $f(\xi, \tau) = G(\xi) h(\tau)$: the weighting function is separable in time and frequency, which is very useful to adapt it to the shape of the AF-signal terms.

We will end this section by presenting other energy distributions that are members of the Cohen's class.




Figure 4.13: Narrow-band ambiguity function of the previous signal : the AF-signal terms are located around the origin, whereas the AF-interference terms are located away from the origin

4.1.4 Other important energy distributions

The Rihaczek and Margenau-Hill distributions

Another possible definition of a time-frequency energy density is given by the Rihaczek distribution. If we consider the interaction energy between a signal x restricted to an infinitesimal interval δ_T centered on t , and x passed through an infinitesimal bandpass filter δ_B centered on ν , it can be approximated by the following expression :

$$\delta_T \delta_B [x(t) X^*(\nu) e^{-j2\pi\nu t}].$$

This leads us to interpret the quantity

$$R_x(t, \nu) = x(t) X^*(\nu) e^{-j2\pi\nu t},$$

called the *Rihaczek distribution*, as a complex energy density at point (t, ν) . This distribution, which corresponds to the element of the Cohen's class for which $f(\xi, \tau) = e^{j\pi\xi\tau}$, verifies many good properties (1-2, 4-11, see section 4.1.1). However, it is complex valued, which can be awkward in practice. It is implemented under the name `tfrrim`. The real part of the Rihaczek distribution is also a time-frequency distribution of the Cohen's class

($f(\xi, \tau) = \cos(\pi\xi\tau)$), known as the *Margenau-Hill distribution* (see the M-file `tfrmh.m`). It has also numerous interesting properties: 1-5, 8, 10-11. As for the WVD, we can define smoothed versions of the Rihaczek and Margenau-Hill distributions. The file `tfrpmh.m` computes the pseudo Margenau-Hill distribution.

The interference structure of the Rihaczek and Margenau-Hill distributions is different from the Wigner-Ville one: the interference terms corresponding to two points located on (t_1, ν_1) and (t_2, ν_2) are positioned at the coordinates (t_1, ν_2) and (t_2, ν_1) . This can be seen on the following example (see fig. 4.14):

```
>> sig=atoms(128, [32,0.15,20,1;96,0.32,20,1]);
>> tfrmh(sig);
```




Figure 4.14: Margenau-Hill distribution of 2 atoms : the position of the interferences is quite different from the one obtained with the WVD

Thus, the use of the Rihaczek (or Margenau-Hill) distribution for signals composed of multi-components located at the same position in time or in frequency is not advised, since the interference terms will then be superposed to the signal terms.

The Page distribution

Motivated by the construction of a causal energy density, Page proposed the following distribution (the *Page distribution*):

$$\begin{aligned} P_x(t, \nu) &= \frac{d}{dt} \left\{ \left| \int_{-\infty}^t x(u) e^{-j2\pi\nu u} du \right|^2 \right\} \\ &= 2 \Re \left\{ x(t) \left(\int_{-\infty}^t x(u) e^{-j2\pi\nu u} du \right)^* e^{-j2\pi\nu t} \right\} \end{aligned}$$

It is the derivative of the energy spectral density of the signal considered before time t . It corresponds to the element of the Cohen's class with parameterization function $f(\xi, \tau) = e^{-j\pi\xi|\tau|}$, and verifies the properties 1-5, 7-10 (see section 4.1.1). Actually, it is the only distribution of the Cohen's class which is simultaneously causal, unitary, compatible with modulations, and preserves time-support.

The function `tfrpage.m` computes this distribution. A frequency-smoothed version of the Page distribution, called the *pseudo-Page distribution*, is also available (see the file `tfrppage.m`).

Joint-smoothings of the WVD

The following distributions correspond to particular cases of the Cohen's class for which the parameterization function depends only on the product of the variables τ and ξ :

$$f(\xi, \tau) = \Phi(\tau\xi) \tag{4.9}$$

where Φ is a decreasing function such that $\Phi(0) = 1$ (the Rihaczek and Margenau-Hill distributions are particular elements of this class). A direct consequence of this definition is that the marginal properties will be respected. Besides, since Φ is a decreasing function, f is a low-pass function, and according to (4.8), this parameterization function will reduce the interferences. That is why these distributions are also known as the *Reduced Interference Distributions*.

- The *Choi-Williams distribution*

One natural choice for Phi is to consider a gaussian function :

$$f(\xi, \tau) = \exp \left[-\frac{(\pi\xi\tau)^2}{2\sigma^2} \right].$$

The corresponding distribution,

$$CW_x(t, \nu) = \sqrt{\frac{2}{\pi}} \int \int_{-\infty}^{+\infty} \frac{\sigma}{|\tau|} e^{-2\sigma^2(s-t)^2/\tau^2} x(s+\frac{\tau}{2}) x^*(s-\frac{\tau}{2}) e^{-j2\pi\nu\tau} ds d\tau$$

is the Choi-Williams distribution. Note that when $\sigma \rightarrow +\infty$, we obtain the WVD. Inversely, the smaller σ , the better the reduction of the interferences. This distribution verifies properties 1-5, 10-11, and can be computed with the M-file `tfrcw.m`. The "cross"-shape of the parameterization function of the Choi-Williams distribution implies that the efficiency of this distribution strongly depends on the nature of the analyzed signal. For instance, if the signal is composed of synchronized components in time or in frequency, the Choi-Williams distribution will present strong interferences. This can be observed on the following example: we analyze four gaussian atoms positioned at the corners of a rectangle rotating around the center of the time-frequency plane (see fig. 4.15):

```
>> load movcw4at
>> clf; movie(M,5);
```

When the time/frequency supports of the atoms overlap, some AF-interference terms are not completely attenuated (those present around the axes of the ambiguity plane), and the efficiency of the distribution is quite poor.

- The *Born-Jordan* and *Zhao-Atlas-Marks distributions*

If we impose to the distributions defined by (4.9) the further condition to preserve time- and frequency- supports, the simplest choice for f is then :

$$f(\xi, \tau) = \frac{\sin(\pi\xi\tau)}{\pi\xi\tau}$$

which defines the *Born-Jordan distribution* :

$$BJ_x(t, \nu) = \int_{-\infty}^{+\infty} \frac{1}{|\tau|} \int_{t-|\tau|/2}^{t+|\tau|/2} x(s + \tau/2) x^*(s - \tau/2) ds e^{-j2\pi\nu\tau} d\tau.$$

Properties 1-5, 8, 10-11 are verified by this distribution, and the corresponding M-file of the Time-Frequency Toolbox is `tfrbj.m`.

If we smooth the Born-Jordan distribution along the frequency axis, we obtain the *Zhao-Atlas-Marks distribution*, defined as

$$ZAM_x(t, \nu) = \int_{-\infty}^{+\infty} \left[h(\tau) \int_{t-|\tau|/2}^{t+|\tau|/2} x(s + \tau/2) x^*(s - \tau/2) ds \right] e^{-j2\pi\nu\tau} d\tau.$$




Figure 4.15: Choi-Williams distribution of 4 atoms rotating around the middle of the time-frequency plane : when the time/frequency supports of the atoms overlap, strong interferences appear on the overlap support

This distribution, also known as the *Cone-Shaped Kernel distribution*, validates properties 3-4 and 8 (only for time) (see the M-file `tfrzam.m` for its computation).

Comparison of the parameterization functions

To illustrate the differences between some of the presented distributions, we represent their weighting (parameterization) function in the ambiguity plane, along with the result obtained by applying them on a two-component signal embedded in white gaussian noise : the signal is the sum of two linear FM signals, the first one with a frequency going from 0.05 to 0.15, and the second one from 0.2 to 0.5. The signal to noise ratio is 10 dB.

On the left-hand side of the figures 4.16 and 4.17, the parameterization functions are represented in a schematic way by the bold contour lines (the weighting functions are mainly non-zeros inside these lines), superimposed to the ambiguity function of the signal. The AF-signal terms are in the middle of the ambiguity plane, whereas the AF-interference terms are distant from the center. On the right-hand side, the corresponding time-frequency distributions are represented.

From these plots, we can conclude that the ambiguity plane is very en-




Figure 4.16: Two chirps embedded in a 10 dB white gaussian noise analyzed by different quadratic distributions. On the left-hand side, the parameterization function is represented by a bold contour line, superimposed to the ambiguity function of the signal. The AF-signal terms are in the middle of the ambiguity plane, whereas the AF-interference terms are distant from the center. On the right-hand side, the corresponding time-frequency distribution is represented

lightening with regard to interference reduction in the case of multicomponent signals. On this example, we notice that the smoothed-pseudo-WVD is a particularly convenient and versatile candidate. This is due to the fact that we can adapt independently the time-width and frequency-width of its weighting function. But in the general case, it is interesting to have several distributions at our disposal since each one is well adapted to a certain type of signal. Besides, for a given signal, as a result of the different interference geometries, these distributions offer complementary descriptions of this signal.




Figure 4.17: Two chirps embedded in a 10 dB white gaussian noise analyzed by different quadratic distributions (concluding)

4.1.5 Conclusion

The Cohen's class, which gather all the quadratic time-frequency distributions covariant by shifts in time and in frequency, offers a wide set of powerful tools to analyze non-stationary signals. The basic idea is to devise a joint function of time and frequency that describes the energy density or intensity of a signal simultaneously in time and in frequency. The most important element of this class is probably the Wigner-Ville distribution, which satisfies many desirable properties. Since these distributions are quadratic, they introduce cross-terms in the time-frequency plane which can disturb the readability of the representation. One way to attenuate these interferences is to smooth the distribution in time and in frequency, according to their structure. But the consequence of this is a decrease of the time and frequency resolutions, and more generally a loss of theoretical properties. The general formulation proposed by Cohen is very useful to have a better understanding of the existing solutions, as well as the connection with the ambiguity function.

But there exists other time-frequency energy distributions, which are not elements of the Cohen's class, i.e. which are not covariant by shifts in time or in frequency. This is the case for example of the affine distributions, which are presented in the next chapter.

4.2 The affine class

Important note: in all the following section, we will consider that the signal (often denoted x) is analytic (see section 2.3 for the definition and the computation of the analytic signal).

For this part, more information can be found in [Fla93], [BB92], [Gon93], [GB96], [Ova94] and [RF92].

4.2.1 Axiomatic definition

The affine group

The Cohen's class, as presented in the previous section, is based on the properties of covariance by shifts in time and in frequency. One important element of this class is the Wigner-Ville distribution, noteworthy for its numerous properties.

In order to favor a time-scale approach of the signal, one can also choose to put forward, among these desirable properties, the *covariance by translation in time and dilation*. The corresponding group of transforms, counterpart of the Weyl-Heisenberg group (see section 3.1.1), is the *affine group*, noted A , already introduced in the context of wavelet transform (see section 3.2.2). Its action induced on a signal $x(t)$ is given by

$$x(t) \rightarrow x_{a',b'}(t) = \frac{1}{\sqrt{|a'|}} x\left(\frac{t-b'}{a'}\right),$$

and on its Fourier transform by

$$X(\nu) \rightarrow X_{a',b'}(\nu) = \sqrt{|a'|} e^{-j2\pi\nu b'} X(a'\nu).$$

General expressions

It is possible to show that if a bilinear time-scale distribution $\Omega_x(t, a)$ is covariant to affine transformations, i.e.

$$\Omega_{x_{a',b'}}(t, a) = \Omega_x\left(\frac{t-b'}{a'}, \frac{a}{a'}\right),$$

then, it is necessarily parameterized as

$$\Omega_x(t, a; \Pi) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Pi\left(\frac{s-t}{a}, a\xi\right) W_x(s, \xi) ds d\xi \quad (4.10)$$

where $\Pi(t, \nu)$ is an arbitrary smoothing function. This distribution will also preserve the signal energy provided that

$$\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Pi(t, \nu) dt \frac{d\nu}{|n\nu|} = 1.$$

The set of such representations defines the *affine class*, which is the class of time-frequency energy distributions covariant by translation in time and dilation. From expression (4.10), it is straightforward that the Wigner-Ville distribution is an element of the affine class: if we introduce an arbitrary non-zero frequency ν_0 , and identify the scale with the inverse of the frequency:

$$a = \frac{\nu_0}{\nu},$$

then the WVD corresponds to the element for which

$$\Pi(t, \nu) = \delta(t) \delta(\nu - \nu_0).$$

A consequence of (4.10) is that the choice of an element in the affine class can be reduced to the choice of an affine correlation kernel $\Pi(t, \nu)$. When Π is a two-dimensional low-pass function, it plays the role of an affine smoothing function which tries to reduce the interferences generated by the WVD.

Another equivalent expression for a generic element can be found in terms of ambiguity:

$$\Omega_x(t, a; \Phi) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Phi(a\xi, \tau/a) A_x(\xi, \tau) e^{-j2\pi\xi t} d\xi d\tau, \quad (4.11)$$

where $\Phi(\xi, \tau)$ is the weighting function corresponding to Π :

$$\Phi(\xi, \tau) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Pi(t, \nu) e^{j2\pi(\nu\tau + \xi t)} dt d\nu,$$

and $A_x(\xi, \tau)$ is the narrow-band ambiguity function already defined in section 4.1.3.

Finally, an alternative characterization of the class (4.10) may be given by using the *bi-frequency kernel* $\Psi(\nu, f)$

$$\Omega_x(t, a; \Pi) = \frac{1}{|a|} \int \int_{-\infty}^{+\infty} \Psi(\nu, f) X\left(\frac{f - \frac{\nu}{2}}{a}\right) X^*\left(\frac{f + \frac{\nu}{2}}{a}\right) e^{-j2\pi\nu t/a} d\nu df \quad (4.12)$$

with

$$\Psi(\nu, f) = \int_{-\infty}^{+\infty} \Pi(t, f) e^{-j2\pi\nu t} dt,$$

where $X(\nu)$ is the Fourier transform of $x(t)$. We will take advantage of these different (but equivalent) expressions of the affine class in the following.

Properties

As for the Cohen's class, it can be useful to impose further constraints on the class defined by (4.10), to obtain a sub-class of distributions which validate particular properties (see page 60). We detail here some of the most important ones.

1. *Energy conservation*: by integrating Ω_x all over the time-scale plane, we obtain the energy of x :

$$E_x = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Omega_x(t, a; \Pi) dt \frac{da}{a^2}$$

2. *Marginal properties*: the energy spectral density and the instantaneous power can be obtained as marginal distributions of Ω_x :

$$\begin{aligned} \int_{-\infty}^{+\infty} \Omega_x(t, a; \Pi) dt &= |X(\frac{\nu_0}{a})|^2 \\ \int_{-\infty}^{+\infty} \Omega_x(t, a; \Pi) \frac{da}{a^2} &= |x(t)|^2 \end{aligned}$$

3. *Real-valued*:

$$\Omega_x(t, a; \Pi) \in \mathcal{R}, \quad \forall t, a$$

4. *Time localization*:

$$X(\nu) = \frac{1}{\sqrt{\nu}} e^{-j2\pi\nu t_0} U(\nu) \Rightarrow \Omega_x(t, \frac{\nu_0}{\nu}; \Pi) = \nu \delta(t - t_0) U(\nu)$$

where $U(\nu)$ is the Heaviside step function.

5. *Unitarity*: conservation of the scalar product from the time domain to the time-scale domain (apart from the squared modulus):

$$\left| \int_{-\infty}^{+\infty} x(t) y^*(t) dt \right|^2 = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Omega_x(t, a; \Pi) \Omega_y^*(t, a; \Pi) dt \frac{da}{a^2}$$

6. *Group delay*: we may want to obtain the group delay of x as the first order moment in time of Ω_x :

$$t_x \left(\frac{\nu_0}{a} \right) = \frac{\int_{-\infty}^{+\infty} t \Omega_x(t, a; \Pi) dt}{\int_{-\infty}^{+\infty} \Omega_x(t, a; \Pi) dt}$$

7. *Narrow-band limit*: it can also be desirable that, for narrow-band signals, the affine distribution Ω_x coincides with the Wigner-Ville distribution:

$$\Omega_x(t, a; \Pi) = W_x \left(t, \frac{\nu_0}{a} \right).$$

4.2.2 Some examples

The scalogram

A first example of affine distribution is given by the *scalogram* (see section 3.4.2). Indeed, it is possible to express it as a smoothed version of the WVD :

$$|T_x(t, a; \Psi)|^2 = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} W_x(s, \xi) W_\Psi\left(\frac{s-t}{a}, a\xi\right) ds d\xi. \quad (4.13)$$

Thus, the scalogram corresponds to the distribution of the affine class for which $\Pi(t, \nu) = W_\Psi(t, \nu)$. Expression (4.13), to be compared with expression (4.7), shows that the scalogram is the affine counterpart of the spectrogram. The scalogram validates properties 1. and 3. and is always positive.

To illustrate the importance of the smoothing operated by Π on the WVD of x , let us consider the case of a Morlet wavelet Ψ . If we note δ_T and δ_B the respectively time and frequency widths of the smoothing operated by the spectrogram of window Ψ (δ_T and δ_B are constant values), these widths become variable with the frequency in the case of the scalogram :

$$\delta_T(\nu) = \nu_0 \delta_T / \nu ; \delta_B(\nu) = \nu \delta_B / \nu_0$$

(ν_0 is the central frequency of the wavelet). This result, already made out in the context of the wavelet transform analysis, is a characteristic of any constant-Q analysis (see section 3.2.1) : at a high frequency, since the signal changes rapidly, a short analysis window is sufficient, whereas at a low frequency, a large window is necessary to identify correctly the pulsation of the signal which changes slowly. However, the importance of the joint smoothing operated by the scalogram is still equivalent to the one of the spectrogram :

$$\delta_T(\nu) \delta_B(\nu) = \delta_T \delta_B.$$

Besides, the trade-off between time and frequency resolutions, following from the Heisenberg-Gabor inequality and which applies to the spectrogram, is also valid for the scalogram.

So as to see the effect of this frequency-dependent smoothing, we analyze with the scalogram (Morlet wavelet) a signal composed of two gaussian atoms, one with a low central frequency, and the other one with a high one (see fig. 4.18) :

```
>> sig=atoms(128, [38,0.1,32,1;96,0.35,32,1]);
>> tfrscalo(sig);
```




Figure 4.18: Morlet scalogram of 2 atoms : the time- and frequency- resolutions depend on the frequency (or scale)

By default, the file `tfrscalom` uses an interactive mode in which you have to specify, from the plot of the spectrum, the approximate lower and higher frequency bounds, as well as the number of samples you wish in frequency (you should indicate here a lower frequency lower than 0.05 and a higher frequency greater than 0.4). The result obtained brings to the fore dependency, with regard to the frequency, of the smoothing applied to the WVD, and consequently of the resolutions in time and frequency.

The product kernel distributions

The formal identification "scale=inverse of the frequency" can be extended to other distributions than the WVD. If we consider kernels of the form

$$\Phi(\xi, \tau) = \Phi(\xi\tau) e^{-j2\pi\nu_0\tau}$$

where ν_0 is some nonzero frequency, we then have the following equivalence between the Cohen's class and the affine class :

$$\Omega_x(t, a; \Phi) = C_x \left(t, \frac{\nu_0}{a}; \Phi \right).$$

The corresponding representations, and in particular the Wigner-Ville, Born-Jordan, Rihaczek and Choi-Williams distributions, are elements of the intersection of these two classes.

The affine smoothed pseudo Wigner distribution : separable kernel

One way to overcome the trade-off between time and frequency resolutions of the scalogram is, as for the smoothed-pseudo-WVD, to use a smoothing function which is separable in time and frequency. The resulting distribution is called the *affine smoothed pseudo Wigner distribution* (noted ASPWD), and writes

$$ASPW_x(t, a) = \frac{1}{a} \int \int_{-\infty}^{+\infty} h\left(\frac{\tau}{a}\right) g\left(\frac{s-t}{a}\right) x\left(s + \frac{\tau}{2}\right) x^*\left(s - \frac{\tau}{2}\right) ds d\tau \quad (4.14)$$

It allows a flexible choice of time and scale resolutions in an independent manner through the choice of the windows g and h . Properties 1. and 3. (see page 85) are satisfied by this distribution provided that g is real and h is hermitian.

As for the SPWVD (see section 4.1.2), the ASPWD allows a continuous passage from the scalogram to the WVD, under the condition that the smoothing functions g and h are gaussian. The time-bandwidth product then goes from 1 (scalogram) to 0 (WVD), with an independent control of the time and frequency resolutions. This is illustrated by the function `movsc2wv.m`, which considers different transitions, on a signal composed of four atoms. To visualize these snapshots, load the mat-file `movsc2wv` and run `movie` (see fig. 4.19):

```
>> load movsc2wv
>> clf; movie(M,10);
```

Here again, the WVD gives the best resolutions (in time and in frequency), but presents the most important interferences, whereas the scalogram gives the worst resolutions, but with nearly no interferences; and the ASPWD allows to choose the best compromise between these two extremes.

To summarize, we have seen that on one hand, the spectrogram is a time-frequency distribution obtained from the WVD by smoothing, and that on the other hand, the scalogram is a time-frequency distribution obtained from the WVD by affine smoothing. The WVD is therefore at the intersection of both classes of time-frequency and time-scale distributions. Besides, it is possible to construct a continuous transition from the spectrogram to the scalogram via the WVD, by changing the smoothing function Π acting on the WVD. The equivalent area of such function Π will vary from zero (we then obtain the "unsmoothed" WVD) to a limit fixed by the Heisenberg-Gabor uncertainty principle (spectrogram and scalogram). This choice corresponds to using the SPWVD's or the ASPWD's with gaussian smoothing functions.




Figure 4.19: Different transitions between the scalogram and the WVD thanks to the ASPWD. The analyzed signal is composed of 4 gaussian atoms

The time-bandwidth product then runs from 0 (WVD) to 1 (spectrogram or scalogram) and truly controls both transitions.

Figure 4.20 illustrates different transitions between the spectrogram and the scalogram on a synthetic signal composed of three gaussian atoms, for different values of BT .

This analysis brings us to the conclusion that, instead of looking at the two extreme representations (spectrogram and scalogram) separately, a deeper insight can be gained by considering a whole continuum between the two extremes, with the WVD as a necessary intermediate step. Moreover, the transition allows a trade-off between joint resolutions and interferences reduction.

The localized bi-frequency kernel distributions

A useful subclass of the affine class consists in characterization functions which are perfectly localized on some curve $f = H(\nu)$ in their bi-frequency




Figure 4.20: From the spectrogram to the scalogram via the WVD

representation (see (4.12)):

$$\Psi(nu, f) = G(\nu) \delta(f - H(\nu)) \Leftrightarrow \Phi(\nu, \tau) = G(\nu) e^{j2\pi H(\nu)\tau}$$

where $G(\nu)$ is an arbitrary function. The corresponding time-scale distributions, which are referred to as *localized bi-frequency kernel distributions*, then read

$$\Omega_x(t, a; \Pi) = \frac{1}{|a|} \int_{-\infty}^{+\infty} G(\nu) X\left(\frac{H(\nu) - \nu/2}{a}\right) X^*\left(\frac{H(\nu) + \nu/2}{a}\right) e^{-j2\pi\nu t/a} d\nu.$$

Actually, it has been shown that the only group delay laws on which a localized bi-frequency kernel distribution can be perfectly localized are power laws (i.e. $t_x(\nu) = t_0 + c\nu^{k-1}$) and logarithmic laws (i.e. $t_x(\nu) = t_0 + c \log \nu$).

As for the product-kernel distributions, with the formal identification $a = \nu_0/\nu$, we can associate to every time-scale distribution of that kind a

time-frequency distribution according to

$$C_x(t, \nu; \Phi) = \Omega_x(t, \nu_0/\nu; \Phi).$$

We give in the following particular examples of such distributions.

- *Bertrand distribution*

If we further impose to these distributions the *a priori* requirements of time localization and unitarity, we obtain

$$G(\nu) = \frac{\nu/2}{\sinh\left(\frac{\nu}{2}\right)} \quad \text{and} \quad H(\nu) = \frac{\nu}{2} \coth\left(\frac{\nu}{2}\right),$$

which leads to the *Bertrand distribution*, defined as

$$B_x(t, a) = \frac{1}{|a|} \int_{-\infty}^{+\infty} \frac{\nu/2}{\sinh\left(\frac{\nu}{2}\right)} X\left(\frac{\nu e^{-\nu/2}}{2a \sinh\left(\frac{\nu}{2}\right)}\right) \times X^*\left(\frac{\nu e^{+\nu/2}}{2a \sinh\left(\frac{\nu}{2}\right)}\right) e^{-j2\pi\nu t/a} d\nu \quad (4.15)$$

It validates properties 1. to 7., except the time-marginal property (see page 85). Besides, we can show that this distribution is the only localized bi-frequency kernel distribution which localizes perfectly the hyperbolic group delay signals:

$$X(\nu) = \frac{e^{j\Phi_x(\nu)}}{\sqrt{\nu}} U(\nu)$$

with $\Phi_x(\nu) = -2\pi \left[\nu t_0 + \alpha \log \frac{\nu}{\nu_c} \right] \Rightarrow B_x(t, a = \frac{\nu_0}{\nu}) = \nu \delta(t - t_x(\nu)) U(\nu)$

where $t_x(\nu) = -\frac{1}{2\pi} \frac{d\Phi_x(\nu)}{d\nu}$ is the group delay. To illustrate this property, consider the signal obtained using the file `gdpower.m` (taken for $k = 0$), and analyze it with the file `tfrbert.m` (see fig. 4.21):

```
>> sig=gdpower(128);
>> tfrbert(sig,1:128,0.01,0.22,128,1);
```

Note that the distribution obtained is well localized on the hyperbolic group delay, but not perfectly: this comes from the fact that the file `tfrbert.m` works only on a subpart of the spectrum, between two bounds f_{min} and f_{max} . Note that the larger the frequency bandwidth, the more needed samples, and consequently the longer the computation time.




Figure 4.21: Bertrand distribution of an hyperbolic group delay signal

- *D-Flandrin distribution*

If we now look for a localized bi-frequency kernel distribution which is real, localized in time and which validates the time-marginal property, we obtain

$$G(\nu) = 1 - (\nu/4)^2 \quad \text{and} \quad H(\nu) = 1 + (\nu/4)^2.$$

The corresponding distribution then writes:

$$D_x(t, a) = \frac{1}{|a|} \int_{-\infty}^{+\infty} (1 - (\nu/4)^2) X \left(\frac{[1 - \nu/4]^2}{a} \right) \times X^* \left(\frac{[1 + \nu/4]^2}{a} \right) e^{-j2\pi\nu t/a} d\nu,$$

which defines the *D-Flandrin distribution*. It validates properties 1-4., 6. and 7. (see page 85), and is the only localized bi-frequency kernel distribution which localizes perfectly signals having a group delay in $\frac{1}{\sqrt{\nu}}$:

$$X(\nu) = \frac{e^{j\Phi_x(\nu)}}{\sqrt{\nu}} U(\nu)$$

with

$$\Phi_x(\nu) = -2\pi[\nu t_0 + 2\alpha\sqrt{\nu}] \Rightarrow D_x \left(t, a = \frac{\nu_0}{\nu} \right) = \nu \delta(t - t_x(\nu)) U(\nu).$$

This can be illustrated using the files `gdpower.m` with $k = 1/2$ and `tfrdfla.m`, as following (see fig. 4.22):

```
>> sig=gdpower(128,1/2);
>> tfrdfla(sig,1:128,0.01,0.22,128,1);
```




Figure 4.22: D-Flandrin distribution of a signal with a group delay in $1/\nu^{1/2}$

Here again, the distribution is almost perfectly localized.

- *Unterberger distributions*

Finally, the choice of

$$G(\nu) = 1 \quad \text{and} \quad H(\nu) = \sqrt{1 + \left(\frac{\nu}{2}\right)^2}$$

corresponds to the *active Unterberger distribution*:

$$U_x^{(a)}(t, a) = \frac{1}{|a|} \int_0^{+\infty} \left(1 + \frac{1}{\alpha^2}\right) X\left(\frac{\alpha}{a}\right) X^*\left(\frac{1}{\alpha a}\right) e^{j2\pi(\alpha-1/\alpha)\frac{t}{a}} d\alpha,$$

which verifies properties 1-4., 6-7. (see page 85) except the time-marginal; and the choice of

$$G(\nu) = \frac{1}{\sqrt{1 + \left(\frac{\nu}{2}\right)^2}} \quad \text{and} \quad H(\nu) = \sqrt{1 + \left(\frac{\nu}{2}\right)^2}$$

corresponds to the *passive Unterberger distribution*:

$$U_x^{(p)}(t, a) = \frac{1}{|a|} \int_0^{+\infty} \frac{2}{\alpha} X\left(\frac{\alpha}{a}\right) X^*\left(\frac{1}{\alpha a}\right) e^{j2\pi(\alpha - \frac{1}{\alpha})\frac{t}{a}} d\alpha,$$

which verifies properties 1-3., 6-7. The active Unterberger distribution is the only localized bi-frequency kernel distribution which localizes perfectly signals having a group delay in $1/\nu^2$:

$$X(\nu) = \frac{e^{j\Phi_x(\nu)}}{\sqrt{\nu}} U(\nu)$$

with

$$\Phi_x(\nu) = -2\pi[\nu t_0 - \alpha/\nu] \Rightarrow U_x^{(a)}(t, a = \nu_0/\nu) = \nu \delta(t - t_x(\nu)) U(\nu).$$

The files `gdpower.m`, considered for $k = -1$, and `tfrunter.m` give us (see fig. 4.23):

```
>> sig=gdpower(128,-1);
>> tfrunter(sig,1:128,'A',0.01,0.22,172,1);
```




Figure 4.23: Active Unterberger distribution of a signal with a group delay in $1/\nu^2$

We will go back over these distributions later on (sub-section 4.2.4) in a different context.

4.2.3 Relation with the ambiguity domain

Need of specific tools for broad-band signals

The WVD, as we have seen in the previous chapter, is a very satisfactory distribution when applied to narrow-band signals. Its use for the description of broad-band signals is also possible, but can lead to surprising images. For example, for an analytic signal whose real part is $\delta(t - t_0)$, the WVD equals to

$$W(t, \nu) = 4 \frac{\sin(4\pi\nu(t - t_0))}{\pi(t - t_0)} U(\nu)$$

where $U(\nu)$ is the Heaviside function, and thus is not well localized in the neighborhood of $t = t_0$ (see fig. 4.24) :

```
>> sig=anapulse(128);  
>> tfrwv(sig);
```




Figure 4.24: WVD of a Dirac impulse at time $t = 64$

Actually, the group of translations in time and frequency (the Weyl-Heisenberg group, see section 3.1.1) on which the WVD is based, and more generally all the Cohen's class, is responsible for these bad localization properties on broad-band signals: since the use of the analytic signal is admitted, the translation in frequency of broad-band signals fails to preserve the frequency

support of the signal (the support of its Fourier transform can not be limited to the positive frequency axis). This suggests to replace the WVD by a distribution more fundamentally based on the affine group.

The *Doppler effect*, which is an important physical phenomenon, provides an additional motivation to use specific methods based on the affine group to analyze broad-band signals. Indeed, it characterizes the fact that a signal returned by a moving target is dilated (or compressed) and delayed compared to the emitted signal. If, for narrow-band emitted signals and low-speed targets (compared to the sound speed in the medium) this phenomenon can be approximated by a translation in time and frequency, for broad-band signals, the dilation of the spectrum has to be taken into account. This is particularly the case in radar and sonar problems where the time-bandwidth product of the emitted signal is important and where the speed of the moving target is often not negligible compared to the wave speed in the medium.

From the Fourier transform to the Mellin transform

A second argument encourages one to find more specific tools to analyze broad-band signals: the eigenvectors of the Weyl-Heisenberg group are the familiar complex exponentials, on which the Fourier transform decomposes a signal, whereas for the affine group, the eigenvectors are hyperbolas. From a slightly different point of view, the Fourier transform is invariant in modulus to translations in frequency, but not to dilations. Therefore, the Fourier transform is no longer the appropriate transform to change the representation space of these signals. It has to be replaced by a new transform, the *Mellin transform*, which is invariant in modulus to dilations, and decomposes the signal on a basis of hyperbolic signals. This transform can be defined as:

$$M_X(\beta) = \int_0^{+\infty} X(\nu) \nu^{j2\pi\beta-1} d\nu$$

where $X(\nu)$ is the Fourier transform of the analytic signal corresponding to $x(t)$. We can show easily that

$$Y(\nu) = X(a\nu) \Rightarrow M_Y(\beta) = a^{-j2\pi\beta} M_X(\beta),$$

which demonstrates the invariance by dilation. The basic elements are waves of the form $\nu^{-j2\pi\beta}$, whose group delay is hyperbolic:

$$t_x(\nu) = \frac{\beta}{\nu}.$$

Thus, the β -parameter can be interpreted as a *hyperbolic modulation rate*, and has no dimension; it is called the *Mellin's scale*.

In the discrete case, the Mellin transform can be calculated rapidly using a fast Fourier transform. Its algorithm, called the *fast Mellin transform*, is computed thanks to the file `fnt.m`. For further details on this transform, see for example [Ova94]. This transform is often used in the Time-Frequency Toolbox to implement functions which are connected to the affine class.

From the narrow-band AF to the wide-band AF

When the signal under analysis can not be considered as narrow-band (i.e. when its bandwidth B is not negligible compared to its central frequency ν_0), the narrow-band ambiguity function is no longer appropriate since the Doppler effect can not be approximated as a frequency-shift. We then consider a *wide-band ambiguity function* (WAF), which can be defined as :

$$\Xi_x(a, \tau) = \frac{1}{\sqrt{a}} \int_{-\infty}^{+\infty} x(t) x^*(t/a - \tau) dt = \sqrt{a} \int_{-\infty}^{+\infty} X(\nu) X^*(a\nu) e^{j2\pi a\tau\nu} d\nu.$$

It corresponds to the wavelet transform of the signal x , whose mother wavelet is the signal x itself. It is then an affine correlation function, which measures the similarity between the signal and its translated (in time) and dilated versions. This ambiguity function can be easily calculated using two Mellin transforms. The M-file `ambifuwb.m` of the Time-Frequency Toolbox computes this expression of the wide-band ambiguity function. To see how it behaves on a practical example, let us consider an Altes signal (see the M-file `altes.m`) (see fig. 4.25) :

```
>> sig=altes(128,0.1,0.45);
>> ambifuwb(sig);
```

The WAF is maximum at the origin of the ambiguity plane.

We can also introduce a symmetric form of the WAF :

$$\Xi_x^{(s)}(\alpha, \tau) = \sqrt{1 - \alpha^2/4} \int_{-\infty}^{+\infty} x\left(\left(1 + \alpha/2\right)t + \frac{\tau}{2}\right) x^*\left(\left(1 - \alpha/2\right)t - \frac{\tau}{2}\right) dt$$

where $a = (1 + \alpha/2)(1 - \alpha/2)$. This expression can be related to the WVD by the following formula :

$$\Xi_x^{(s)}(\alpha, \tau) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \sqrt{1 - \alpha^2/4} e^{j2\pi(\tau+\alpha t)\nu} W_x(t, \nu) d\nu.$$




Figure 4.25: Wide-band ambiguity function of an Altes signal

From the WVD to the Bertrand distribution

Now that we defined the symmetric wide-band ambiguity function, it would be interesting to obtain an expression equivalent to the one linking the WVD and the narrow-band ambiguity function, but replaced in the affine context. This can be done by applying a Fourier transform to the delay variable of the symmetric WAF, and a Mellin transform to the α variable :

$$\begin{aligned}
 B_x(t, \nu) &= \int_0^{+\infty} \int_{-\infty}^{+\infty} \Xi_x^{(s)}(\alpha, \tau) e^{-j2\pi\nu\tau} \alpha^{j2\pi t-1} d\tau d\alpha \\
 &= \nu \int_{-\infty}^{+\infty} \frac{u/2}{\sinh\left(\frac{u}{2}\right)} X\left(\frac{\nu u e^{-u/2}}{2 \sinh\left(\frac{u}{2}\right)}\right) X^*\left(\frac{\nu u e^{+u/2}}{2 \sinh\left(\frac{u}{2}\right)}\right) e^{-j2\pi\nu u t} du \quad (4.16)
 \end{aligned}$$

which corresponds to the Bertrand distribution, already introduced in section 4.2.2 (the equivalence between formula (4.15) and (4.16) is obtained by identifying ν as the inverse of the scale: $\nu = \frac{\nu_0}{a}$ with $\nu_0 = 1$ Hz).

4.2.4 The affine Wigner distributions

Introduction

The Bertrand distribution B_x given by (4.15) or (4.16) is in fact covariant by a larger group than the affine group A : this group, G_0 , of transformations

$g = (a, b, c)$, where (a, b) is an element of A and c is real, acts on the signal X as:

$$X(\nu) \rightarrow X_g(\nu) = \sqrt{|a|} e^{-j2\pi(\nu b + c \ln(\nu))} X(a\nu).$$

The resulting change on B_x is:

$$B_x \rightarrow B_x^g(t, \nu) = B_x \left(\frac{t - b - c/\nu}{a}, a\nu \right).$$

Actually, it is possible to generalize this extended covariance property to a sub-class of affine distributions, not only restricted to the Bertrand distribution. It can be shown that the only three-parameter groups, noted G_k , including the affine group, are defined as follows: for $k \neq 1$, G_k is the group of elements $g = (a, b, c)$ with composition law:

$$gg' = (aa', b + ab', c + a^k c').$$

Group G_1 has a slightly different composition law:

$$gg' = (aa', b + ab' + a \ln(a)c', c + ac').$$

The action of these groups on the analytic signal $X(\nu)$ is then dependent on k according to:

$$\begin{aligned} X(\nu) \rightarrow X_g(\nu) &= \sqrt{|a|} X(a\nu) e^{-j2\pi(\nu b + c\nu^k)} && \text{for } k \neq 0, 1; \\ X(\nu) \rightarrow X_g(\nu) &= \sqrt{|a|} X(a\nu) e^{-j2\pi(\nu b + c \ln(\nu))} && \text{for } k = 0; \\ X(\nu) \rightarrow X_g(\nu) &= \sqrt{|a|} X(a\nu) e^{-j2\pi(\nu b + c\nu \ln(\nu))} && \text{for } k = 1. \end{aligned}$$

The distributions P_x^k covariant by these three-parameter solvable groups G_k , and satisfying the time-reversal invariance ($Y(\nu) = X^*(\nu) \Rightarrow P_y^k(t, \nu) = P_x^k(-t, \nu)$), are then found to be:

$$P_x^k(t, \nu) = \int_{-\infty}^{+\infty} \nu \mu_k(u) X(\lambda_k(u)\nu) X^*(\lambda_k(-u)\nu) e^{j2\pi(\lambda_k(u) - \lambda_k(-u))t\nu} du \quad (4.17)$$

$$\text{where } \lambda_k(u) = \left(\frac{k(e^{-u} - 1)}{e^{-ku} - 1} \right)^{\frac{1}{k-1}}$$

and $\mu_k(u)$ is a real positive and even function. The definition (4.17) is valid for any real k provided that $\lambda_k(u)$ is defined by continuity for $k = 0$ and $k = 1$:

$$\lambda_0(u) = -\frac{u}{e^{-u} - 1} \quad \text{and} \quad \lambda_1(u) = \exp \left(1 + \frac{ue^{-u}}{e^{-u} - 1} \right).$$

Expression (4.17) defines the *class of affine Wigner distributions*. As we will see in the next section, this class, introduced on mathematical considerations, is equivalent to the class of localized bi-frequency kernel distributions (see section 4.2.2). We now investigate special cases of μ_k leading to distributions satisfying unitarity and/or localization properties.

Some examples

Two special families of affine Wigner distributions can be determined by imposing constraints on P_x^k . The first one is unitarity (see page 85, property 5.), which is satisfied if μ_k is given by

$$\mu_k(u) = \sqrt{\lambda_k(u) \lambda_k(-u)} \frac{d(\lambda_k(u) - \lambda_k(-u))}{du}.$$

The second one is time-localization (property 4.), which implies that

$$\mu_k(u) = \sqrt{\lambda_k(u) \lambda_k(-u)} \frac{d(\lambda_k(u) - \lambda_k(-u))}{du}.$$

- $k = 0$: *the Bertrand distribution*

The choice of $k = 0$ under one or the other (or both) constraints leads to the Bertrand distribution, already defined in sections 4.2.2 and 4.2.3: $P_x^0(t, \nu) = B_x(t, \nu)$. In fact, it is the only affine Wigner distribution which satisfies simultaneously the unitarity and the time localization.

- $k = 2$: *the Wigner-Ville distribution*

The unitary affine Wigner distribution corresponding to $k = 2$ is the Wigner-Ville distribution (see section 4.1.1) provided that x is analytic: $P_x^2(t, \nu) = W_x(t, \nu)$.

- $k = 1/2$: *The D-Flandrin distribution*

The time-localization constraint together with the choice $k = 1/2$ leads to the D-Flandrin distribution, already defined in section 4.2.2: $P_x^{1/2}(t, \nu) = D_x(t, \nu)$.

- $k = -1$: *The active Unterberger distribution*

Another known example of time-localized distribution is obtained for $k = -1$: it corresponds to the active Unterberger distribution (see section 4.2.2). While this form is non-unitary, it cooperates with its passive form to produce an isometry-like relation:

$$\int_{-\infty}^{+\infty} \int_0^{+\infty} U_x^{(a)}(t, \nu) U_y^{(p)*}(t, \nu) d\nu dt = \left| \int_{-\infty}^{+\infty} x(u) y^*(u) du \right|^2.$$

- $k \rightarrow \pm\infty$: *The Margenau-Hill distribution*

Finally, under the unitarity constraint, it is interesting to consider the two distributions obtained for $k \rightarrow -\infty$ and $k \rightarrow +\infty$: if we note respectively P_- and P_+ these two distributions and take their arithmetic mean, we obtain exactly the Margenau-Hill distribution (see section 4.1.4):

$$\frac{P_x^+(t, \nu) + P_x^-(t, \nu)}{2} = \Re \{R_x(t, \nu)\}.$$

Interference structure

The interference structure of the affine Wigner distributions can be determined thanks to the following geometric argument: two points (t_1, ν_1) and (t_2, ν_2) belonging to the trajectory on which a distribution is localized interfere on a third point (t_i, ν_i) which is necessarily located on the same trajectory. Consequently, using the result of section 4.2.2 which says that the localized bi-frequency kernel distributions are localized on power law group delays of the form $t_x(\nu) = t_0 + c\nu^{k-1}$, one can show that the coordinates (t_i, ν_i) are determined by the relation ([GF92])

$$\omega_i = \left(\frac{\omega_2^k - \omega_1^k}{k(\omega_2 - \omega_1)} \right)^{\frac{1}{k-1}}$$

where $\omega = \nu$ or $\omega = (t - t_0)^{\frac{1}{k-1}}$. These "mid-point" coordinates can be computed using the M-file `midpoint.m` of the Time-Frequency Toolbox. Figure 4.26 represents the location of interference point corresponding to two points of the time-frequency plane (t_1, f_1) and (t_2, f_2) , for different values of k . In particular, for $k = 2$, corresponding to the Wigner-Ville distribution, we obtain the geometric mid-point.

To illustrate this interference geometry, let us consider the case of a signal with a sinusoidal frequency modulation:

```
>> [sig,ifl]=fmsin(128);
```

The file `plotsid.m` allows one to construct the interferences of an affine Wigner distribution perfectly localized on a power-law group-delay (specifying k), for a given instantaneous frequency law (or the superposition of different instantaneous frequency laws). For example, if we consider the case of the Bertrand distribution ($k = 0$) (see fig. 4.27),

```
>> plotsid(1:128,ifl,0);
```




Figure 4.26: Locus of the interferences between 2 points for the affine Wigner distributions (parameterized by k). For $k = 2$, which corresponds to the Wigner-Ville distribution, we obtain the geometric mid-point

we obtain an interference structure completely different from the one obtained for the Wigner-Ville distribution ($k = 2$) (see fig. 4.28) :

```
>> plotsid(1:128,if1,2);
```

For the active Unterberger distribution ($k = -1$), the result is the following (see fig. 4.29) :

```
>> plotsid(1:128,if1,-1);
```

We can notice the presence of an inflexion point (corresponding to the intersection of an infinite number of lines joining two symmetric points from the sinusoid) in the case of the WVD distribution, which disappears in the other distributions.

4.2.5 The pseudo affine Wigner distributions

The affine Wigner distributions (4.17) show great potential as flexible tools for time-varying spectral analysis. However, as for some distributions of the Cohen's class, they present two major practical limitations : first the entire signal enters into the calculation of these distributions at every point




Figure 4.27: Theoretical diagram of the interferences of the Bertrand distribution for a sinusoidal frequency modulation

(t, ν) , and second, due to their nonlinearity, interference components arise between each pair of signal components. To overcome these limitations, a set of (smoothed) pseudo affine Wigner distributions has been introduced recently. We present here the main results relative to this new class of affine distributions.

Derivation

Recall from section 4.1.1 that we obtained the pseudo Wigner-Ville distribution by introducing a window function into the Wigner-Ville distribution. An analogous windowing operated on the affine Wigner distributions (4.17) leads to the *pseudo affine Wigner distributions*. But in contrast to the pseudo Wigner-Ville case, this windowing must be frequency-dependent, to ensure that the resulting time-scale distribution remains scale-covariant. As a result, the smoothing in frequency is constant-Q, rather than constant-bandwidth as in the pseudo Wigner-Ville distribution. The general expression of this new class of distributions, expressed in the time-domain, writes:

$$\begin{aligned} \tilde{P}_x^k(t, \nu) = & \nu \int_{-\infty}^{+\infty} \mu_k(u) \left[\int_{-\infty}^{+\infty} x(\tau) h[\nu\lambda_k(u)(\tau - t)] e^{-j2\pi\lambda_k(u)\nu(\tau-t)} d\tau \right] \\ & \times \left[\int_{-\infty}^{+\infty} x(\tau_p) h[\nu\lambda_k(-u)(\tau_p - t)] e^{-j2\pi\lambda_k(-u)\nu(\tau_p-t)} d\tau_p \right]^* du \end{aligned} \quad (4.18)$$




Figure 4.28: Theoretical diagram of the interferences of the Wigner-Ville distribution for a sinusoidal frequency modulation

where h is the time-windowing function. By analogy with the pseudo Wigner-Ville distributions, we call these distributions the pseudo affine Wigner distributions.

An efficient online implementation can be obtained if we reorder (4.18) to yield

$$\tilde{P}_x^k(t, \nu) = \int_{-\infty}^{+\infty} \frac{\mu_k(u)}{\sqrt{\lambda_k(u)\lambda_k(-u)}} T_x(t, \lambda_k(u)\nu; \Psi) T_x^*(t, \lambda_k(-u)\nu; \Psi) du \quad (4.19)$$

where $T_x(t, \nu; \Psi)$ is the continuous wavelet transform (see section 3.2.1), and $\Psi(\tau) = h(\tau) e^{j2\pi\tau}$ is a bandpass wavelet function.

Time-frequency smoothing interpretation

The time-windowing function h introduced in (4.18) or (4.19) attenuates interference components that oscillate in the frequency direction. To suppress interference terms oscillating in the time direction, we must smooth in that direction with a low-pass function G . The resulting distributions

$$\tilde{P}_x^k(t, \nu) = \int_{-\infty}^{+\infty} G(u) \frac{\mu_k(u)}{\sqrt{\lambda_k(u)\lambda_k(-u)}} T_x(t, \lambda_k(u)\nu; \Psi) \times T_x^*(t, \lambda_k(-u)\nu; \Psi) du, \quad (4.20)$$




Figure 4.29: Theoretical diagram of the interferences of the active Unterberger distribution for a sinusoidal frequency modulation

are called the *smoothed pseudo affine Wigner distributions*. It is important to notice that, like the (smoothed) pseudo Wigner-Ville case with the localization on linear chirps, (smoothed) pseudo affine Wigner distributions are no longer localized on power-law group delays. Nevertheless, as Q (the quality factor of the wavelet Ψ) tends towards infinity and $G(u)$ to the all-pass function, this localization property is asymptotically recovered since \tilde{P}_x^k converges to P_x^k . Besides, since (4.20) can be implemented efficiently, this convergence property provides us with an efficient-implementation approximation of any affine Wigner distribution (by considering the corresponding pseudo affine Wigner distribution with a large Q).

Expression (4.20) is used in the function `tfrspaw.m` which computes these (smoothed) pseudo affine Wigner distributions.

Examples

Finally, we present two examples of such distributions for different values of k , and analyze the results obtained on a real echolocation signal from a bat. This signal is obtained from the file `bat.mat` :

```
>> load bat; N=128;
>> sig=hilbert(bat(801:7:800+N*7)');
```

For each value of k , we compute the corresponding affine Wigner distribution and smoothed pseudo affine Wigner distribution.

- $k = 2$: *affine smoothed pseudo Wigner distribution*

In this case, (4.20) becomes the affine smoothed pseudo Wigner distribution, already introduced in section 4.2.2 on separable kernel considerations.

```
>> tfrwv(sig);
>> tfrspaw(sig,1:N,2,24,0,0.1,0.4,N,1);
```




Figure 4.30: WVD of a bat sonar signal

On figure 4.30, the WVD presents interference terms because of the non-linearity of the frequency modulation, whereas on figure 4.31, the affine frequency smoothing operated by the affine smoothed pseudo Wigner distribution almost perfectly suppresses the interference terms.

- $k = 0$: *pseudo Bertrand distribution*

This value of k reduces (4.20) to a simple expression

$$\tilde{P}_x^0(t, \nu) = \int_{-\infty}^{+\infty} G(u) T_x(t, \lambda_k(u)\nu; \Psi) T_x^*(t, \lambda_k(-u)\nu; \Psi) du$$

which is called the (smoothed) pseudo Bertrand distribution.




Figure 4.31: Affine smoothed pseudo Wigner distribution of the bat sonar signal

```
>> tfrbert(sig,1:N,0.1,0.4,N,1);
>> tfrspaw(sig,1:N,0,32,0,0.1,0.4,N,1);
```

Figure 4.32 represents the Bertrand distribution. The approximate hyperbolic group delay law of the bat signal explains the good result obtained with this distribution (compared to the WVD). However, it remains some interference terms, which are almost perfectly cancelled on figure 4.33 (pseudo Bertrand distribution).

4.2.6 Conclusion

The constraint of affine covariance has been shown in this part to be relevant for the derivation of time-frequency representations. It leads to a class of affine distributions which is the counterpart of the Cohen's class associated to time and frequency translations. These two classes can also be seen as a result of some 2D correlation acting on the WVD. We have thereby derived a large class of time-scale and time-frequency representations, on which many possible (and sometimes exclusive) properties may be imposed. We have studied several specific requirements (such as energy normalization, time marginals ...) and associated parameterization of the representation. There is obviously a great versatility for the choice of representations, which




Figure 4.32: Bertrand distribution of the bat sonar signal

may be appropriate for various applications. Each one is appropriate to describe properly specific characteristics of a signal, and one has to benefit from the complementarity of these tools. Vice versa, a good interpretation of the time-frequency and time-scale images necessitates a deep knowledge of the mechanisms of information's allocation in the plane.

4.3 The reassignment method

4.3.1 Introduction

Bilinear time-frequency distributions, presented in the previous two sections, offer a wide range of methods designed for the analyze of non stationary signals. Nevertheless, a critical point of these methods is their readability, which means both a good concentration of the signal components and no misleading interference terms. Some efforts have been made recently in that direction, and in particular a general methodology referred to as *reassignment*. The purpose of this section is to present this methodology, to illustrate it on different examples, and to make the link with connected approaches (see [AF94], [KdVG76] and [AF95] for more details on reassignment).




Figure 4.33: Pseudo-Bertrand distribution of the bat sonar signal

4.3.2 The reassignment of the spectrogram

The original idea of reassignment was introduced in an attempt to improve the spectrogram. Indeed, as any other bilinear energy distribution, the spectrogram is faced with an unavoidable trade-off between the reduction of misleading interference terms and a sharp localization of the signal components.

Let us recall the expression of the spectrogram as a 2D-convolution of the Wigner-Ville distribution of the signal by the WVD of the analysis window :

$$S_x(t, \nu; h) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} W_x(s, \xi) W_h(t - s, \nu - \xi) ds d\xi. \quad (4.21)$$

Therefore, this distribution reduces the interference terms of the signal's WVD, but at the expense of opposed time and frequency resolutions, and of biased marginals and first order moments. However, a closer look at expression (4.21) shows that $W_h(t - s, \nu - \xi)$ delimits a time-frequency domain at the vicinity of the (t, ν) point, inside which a weighted average of the signal's WVD values is performed. The key point of the reassignment principle is that these values have no reason to be symmetrically distributed around (t, ν) , which is the geometrical center of this domain. Therefore, their average should not be assigned at this point, but rather at the center of gravity of this domain, which is much more representative of the local energetic distribution of the signal. Reasoning with a mechanical analogy, the local energy

distribution $W_h(t-s, \nu-\xi)W_x(s, \xi)$ (as a function of s and ξ) can be considered as a mass distribution, and it is much more accurate to assign the total mass (i.e. the spectrogram value) to the center of gravity of the domain rather than to its geometrical center.

This is exactly how the reassignment method proceeds : it moves each value of the spectrogram computed at any point (t, ν) to another point $(\hat{t}, \hat{\nu})$ which is the center of gravity of the signal energy distribution around (t, ν) :

$$\hat{t}(x; t, \nu) = \frac{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} s W_h(t-s, \nu-\xi) W_x(s, \xi) ds d\xi}{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} W_h(t-s, \nu-\xi) W_x(s, \xi) ds d\xi} \quad (4.22)$$

$$\hat{\nu}(x; t, \nu) = \frac{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \xi W_h(t-s, \nu-\xi) W_x(s, \xi) ds d\xi}{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} W_h(t-s, \nu-\xi) W_x(s, \xi) ds d\xi} \quad (4.23)$$

and thus leads to a reassigned spectrogram, whose value at any point (t', ν') is the sum of all the spectrogram values reassigned to this point :

$$S_x^{(r)}(t', \nu'; h) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} S_x(t, \nu; h) \delta(t' - \hat{t}(x; t, \nu)) \delta(\nu' - \hat{\nu}(x; t, \nu)) dt d\nu \quad (4.24)$$

One of the mostly interesting properties of this new distribution is that it also uses the phase information of the short-time Fourier transform, and not only its squared modulus as in the spectrogram. This can be seen from the following expressions of the reassignment operators :

$$\begin{aligned} \hat{t}(x; t, \nu) &= -\frac{d\Phi_x(t, \nu; h)}{d\nu} \\ \hat{\nu}(x; t, \nu) &= \nu + \frac{d\Phi_x(t, \nu; h)}{dt} \end{aligned}$$

where $\Phi_x(t, \nu; h)$ is the phase of the STFT of x : $\Phi_x(t, \nu; h) = \arg(F_x(t, \nu; h))$. However, these expressions do not lead to an efficient implementation, and have to be replaced by the following ones :

$$\begin{aligned} \hat{t}(x; t, \nu) &= t - \Re \left\{ \frac{F_x(t, \nu; T_h) F_x^*(t, \nu; h)}{|F_x(t, \nu; h)|^2} \right\} \\ \hat{\nu}(x; t, \nu) &= \nu - \Im \left\{ \frac{F_x(t, \nu; D_h) F_x^*(t, \nu; h)}{|F_x(t, \nu; h)|^2} \right\} \end{aligned}$$

where $T_h(t) = t \times h(t)$ and $D_h(t) = \frac{dh}{dt}(t)$. Reassigned spectrograms are therefore very easy to implement, and do not require a drastic increase in computational complexity.

Finally, it should also be underlined that the reassigned spectrogram, though no longer bilinear, satisfies the time and frequency shifts covariance,

the energy conservation (provided that $h(t)$ is of unit energy), and the non-negativity property. It can also be shown that, since the WVD is perfectly localized on linear chirp signals and impulses, any reassigned spectrogram also satisfies this property :

$$\begin{aligned} x(t) = A \exp \left\{ j \left\{ \nu_0 t + \alpha t^2 / 2 \right\} \right\} &\Rightarrow \hat{\nu} = \nu_0 + \alpha \hat{t} \\ x(t) = A \delta(t - t_0) &\Rightarrow \hat{t} = t_0. \end{aligned}$$

Before presenting the generalization of this method to the Cohen's class and to the affine class, let us have a look at the readability improvement obtained by the reassigned spectrogram on an example of multi-component signal. The reassigned spectrogram is available thanks to the M-file `tfrrsp.m`. The result is compared to the spectrogram and to the "ideal" representation (`tfrideal.m`) based on the knowledge of the instantaneous frequency law of each component :

```
>> N=128; [sig1 ifl1]=fmsin(N,0.15,0.45,100,1,0.4,-1);
>> [sig2 ifl2]=fmhyp(N,[1 .5],[32 0.05]);
>> sig=sig1+sig2;
>> tfrideal([ifl1 ifl2]);
>> figure; tfrrsp(sig);
```

The file `tfrrsp.m` allows you to display the spectrogram itself or its reassigned version. The improvement given by the reassignment method is obvious : the two components are much better localized and almost perfectly concentrated, and there are very few cross-terms.

4.3.3 Reassignment of the Cohen's class representations

The presentation of the reassignment principle done above allows a straightforward extension of its use to other distributions. Indeed, if we consider the general expression of a distribution of the Cohen's class as a 2D-convolution of the WVD,

$$C_x(t, \nu; \Pi) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Pi(t - s, \nu - \xi) W_x(s, \xi) ds d\xi,$$

replacing the particular smoothing kernel $W_h(u, \xi)$ in expressions (4.22), (4.23) and (4.24) by an arbitrary kernel $\Pi(s, \xi)$ simply defines the reassignment of any member of the Cohen's class :

$$\hat{t}(x; t, \nu) = \frac{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} s \Pi(t - s, \nu - \xi) W_x(s, \xi) ds d\xi}{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Pi(t - s, \nu - \xi) W_x(s, \xi) ds d\xi}$$




Figure 4.34: Reassignment of the spectrogram on a synthetic signal composed of a sinusoidal frequency modulation simultaneously with a hyperbolic frequency modulation : comparison with the “ideal” time-frequency representation and with the spectrogram

$$\hat{\nu}(x; t, \nu) = \frac{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \xi \Pi(t-s, \nu-\xi) W_x(s, \xi) ds d\xi}{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Pi(t-s, \nu-\xi) W_x(s, \xi) ds d\xi}$$

$$C_x^{(r)}(t', \nu'; \Pi) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} C_x(t, \nu; \Pi) \delta(t' - \hat{t}(x; t, \nu)) \delta(\nu' - \hat{\nu}(x; t, \nu)) dt d\nu.$$

The resulting reassigned distributions efficiently combine a reduction of the interference terms provided by a well adapted smoothing kernel and an increased concentration of the signal components achieved by the reassignment. From a theoretical point of view, these distributions are covariant by time and frequency shifts, and are perfectly localized for linear chirp signals and impulses. Finally, for the most common cases, such as the SPWVD and the Reduced Interference Distributions (see section 4.1.2 and 4.1.4), the reassignment operators $\hat{t}(x; t, \nu)$ and $\hat{\nu}(x; t, \nu)$ are almost as easy to compute as for the spectrogram.

4.3.4 Reassignment of the affine class representations

Similarly, the reassignment method can also be applied to the time-scale energy distributions. Starting from the general expression :

$$\Omega_x(t, a; \Pi) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Pi(s/a, \nu_0 - a\xi) W_x(t - s, \xi) ds d\xi$$

we can see that the representation value at any point $(t, a = \nu_0/\nu)$ is the average of the weighted WVD values on the points $(t - s, \xi)$ located in a domain centered on (t, ν) and bounded by the essential support of Π . In order to avoid the resultant signal components broadening while preserving the cross-terms attenuation, it seems once again appropriate to assign this average to the center of gravity of these energy measures, whose coordinates are :

$$\begin{aligned} \hat{t}(x; t, \nu) &= t - \frac{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} s \Pi(s/a, \nu_0 - a\xi) W_x(t - s, \xi) ds d\xi}{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Pi(s/a, \nu_0 - a\xi) W_x(t - s, \xi) ds d\xi} \\ \hat{\nu}(x; t, \nu) &= \frac{\nu_0}{\hat{a}(x; t, \nu)} \\ &= \frac{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \xi \Pi(s/a, \nu_0 - a\xi) W_x(t - s, \xi) ds d\xi}{\int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \Pi(s/a, \nu_0 - a\xi) W_x(t - s, \xi) ds d\xi} \end{aligned}$$

rather than to the point $(t, a = \nu_0/\nu)$ where it is computed. The value of the resulting modified time-scale representation on any point (t', a') is then the sum of all the representation values moved to this point :

$$\Omega_x^{(r)}(t', a'; \Pi) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} a'^2 \Omega_x(t, a; \Pi) \delta(t' - \hat{t}(x; t, a)) \delta(a' - \hat{a}(x; t, a)) dt \frac{da}{a^2}.$$

As for the Cohen's class, it can be shown that these modified distributions are no longer bilinear, but are covariant by time shifts and time scalings, distribute the energy of the signal on the whole time-scale plane, and are also perfectly localized for chirps and impulses.

4.3.5 Numerical examples

In order to evaluate the benefits of the reassignment method in practical applications, a comparison of the experimental results provided by some time-frequency representations and their modified versions is shown in this section. The analyzed signal is a 128-points signal made up of a sinusoidal frequency modulation followed by a pure tone simultaneously with a chirp component :

```

>> [sig1 ifl1]=fmsin(60,0.15,0.35,50,1,0.35,1);
>> [sig2 ifl2]=fmlin(60,0.3,0.1);
>> [sig3 ifl3]=fmconst(60,0.4);
>> sig=[sig1;zeros(8,1);sig2+sig3];
>> iflaw=zeros(128,2);
>> iflaw(:,1)=[ifl1;NaN*ones(8,1);ifl2];
>> iflaw(:,2)=[NaN*ones(68,1);ifl3];

```

We first plot the instantaneous frequency laws (obtained by `tfrideal`), to which the proposed solutions should be as near as possible, and the WVD of this signal (see the first two plots of figure 4.35):

```

>> tfrideal(iflaw);
>> figure; tfrwv(sig);

```

With the WVD, the signal components are well localized, but the numerous cross-terms make the figure hardly readable. If we now consider the smoothed pseudo-WVD and its reassigned version (see the third and fourth plots of fig. 4.35),

```

>> tfrrspwv(sig);

```

we can see that the smoothing done by the SPWVD almost completely suppress the cross terms, but the signal components localization becomes coarser. The improvement given by the reassignment method is obvious : all components are much better localized, leading to a nearly ideal representation. The next distributions we consider are the spectrogram (see the first two plots of fig. 4.36) and the Morlet scalogram (see the third and fourth plots of fig. 4.36):

```

>> figure(1); tfrrsp(sig);
>> figure(2); tfrrmsc(sig);

```

These two distributions present nearly no cross terms, except at the bottom of the sinusoid and around time $t = 64$. But the time and frequency resolutions are not good, especially at low frequencies in the case of the scalogram. The reassignment method improves considerably these localizations, and the reassigned spectrogram is even perfectly concentrated for the chirp components. The result obtained with the modified scalogram is less good, especially at low frequencies where the time-resolution is really inadequate.

Finally, we represent the pseudo-Page and the pseudo Margenau-Hill distributions with their reassigned version (see fig. 4.37):

```
>> figure(1); tfrrppag(sig);
>> figure(2); tfrrpmh(sig);
```

These representations (before reassignment) are hardly readable since some cross-terms are superimposed on the signal components. Their modified versions give much better localized signal components, but less concentrated than in the case of the spectrogram or the SPWVD.

4.3.6 Connected approaches

Connections of the reassignment method has been found with other techniques which extract relevant information from the time-frequency plane.

Friedman's instantaneous frequency density

A first example is the instantaneous frequency density : so as to take advantage of the phase structure of the short-time Fourier transform (STFT), Friedman simply computed at each time t the histogram of the frequency displacements $\hat{\nu}(x; t, \nu)$ of the spectrogram. The resulting time-frequency representation is no more an energy distribution, and could be derived as well from any other reassigned distribution.

Here is an example of this instantaneous frequency density, obtained with the M-file `friedman.m` on the pseudo-WVD of the previous signal (see fig. 4.38):

```
>> t=1:2:127; [tfr,rtfr,hat]=tfrrpww(sig,t);
>> friedman(tfr,hat,t,'tfrsp',1);
```

Although some cross terms are still present, the localization of the components is quite good, especially for the chirp components.

Extraction of ridges and skeleton

Another related approach is the extraction of *ridges* and *skeleton*. This method extracts from either the STFT or the continuous wavelet transform (CWT) some particular sets of curves deduced from the stationary points of their phase (see [Fla93] for more information about the stationary phase principle). Indeed, applying the stationary phase theorem to the signal reconstruction formula of the CWT $T_x(t, a; \Psi)$ expressed in the frequency domain :

$$X(\nu) = \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \sqrt{a} H(a\nu) T_x(t, a; \Psi) e^{-j2\pi\nu t} dt \frac{da}{a^2}$$

leads to particular points such that

$$\hat{t}(x; t, a) = t - \Phi'_h(\nu_0) \quad \text{and} \quad \hat{a}(x; t, a) = a, \quad (4.25)$$

with $\Phi_h(\nu) = \arg\{H(\nu)\}$, and which constitute a set of curves called the *horizontal ridges* of the representation.

Similarly, applying the stationary phase principle to the signal reconstruction formula of the CWT expressed in the time domain leads to particular points such that

$$\hat{t}(x; t, a) = t \quad \text{and} \quad \hat{a}(x; t, a) = a \frac{\nu_0}{\phi'_h(0)}, \quad (4.26)$$

with $\phi_h(t) = \arg\{h(t)\}$, and which constitute a set of curves called the *vertical ridges* of the representation. These relations between the ridges and the reassignment operators suggest to extract the ridges of any reassigned distribution by a straightforward generalization of expressions (4.25), (4.26).

For example, let us extract the ridges from the spectrogram of the previous signal (see fig. 4.39):

```
>> [tfr,rtfr,hat]=tfrrsp(sig);
>> ridges(tfr,hat);
```

The result is interesting : apart from some “gaps” present in particular on the sinusoidal frequency modulation, this method concentrates and localizes nearly ideally the signal in the time-frequency plane, even when there are two components present at the same time (or at the same frequency).

4.3.7 Conclusion

The reassignment method creates a modified version of a time-frequency representation by moving the representation values away from where they are computed. These displacements depend on the signal and on the representation, forcing the bilinearity to be lost, but they are still consistent with many of the representation properties. The principle of reassignment exploits the local structures of a distribution in both time and frequency directions. The experimental results show that this method provides a higher concentration in the time-frequency plane, but of course does not remove all the cross terms.




Figure 4.35: Comparison of different time-frequency distributions and their reassigned version (1/3) : the analyzed signal is composed of three components, as can be clearly seen on the first plot representing the instantaneous frequency laws of the components. The other plots are the Wigner-Ville distribution, the smoothed pseudo Wigner-Ville distribution and its reassigned version




Figure 4.36: Comparison of different time-frequency distributions and their reassigned version (2/3) : the spectrogram and the Morlet scalogram




Figure 4.37: Comparison of different time-frequency distributions and their reassigned version (3/3) : the pseudo Page distribution and the pseudo Margenau-Hill distribution




Figure 4.38: Instantaneous frequency density defined by Friedman, computed from the frequency displacements $\hat{\nu}(x; t, \nu)$ of the pseudo-WVD




Figure 4.39: Extraction of ridges from the spectrogram

Chapter 5

Extraction of information from a time-frequency image

Up to this point, we have examined the main solutions proposed to the problem of representing a non-stationary signal in the time-frequency plane. We now consider the problem of the interpretation of the time-frequency image which describes the evolution with time of the frequency content of the signal. Even if they all tend to the same goal, each representation has to be interpreted differently, according to its own properties. For example, some of them present important interference terms, other are only positive, other are perfectly localized on particular signals. . . So the extraction of information has to be done with care, from the knowledge of these properties. We give in the following some general guide lines to profit from a time-frequency image.

5.1 Moments and marginals

The moments and marginals of some representations provide important information about the signal, like its amplitude modulation or its instantaneous frequency, for example with the aim of demodulating the signal.

5.1.1 Moments

The first and second order moments, in time and in frequency, of a time-frequency energy distribution tfr are defined as

$$f_m(t) = \frac{\int_{-\infty}^{+\infty} f \text{tfr}(t, f) df}{\int_{-\infty}^{+\infty} \text{tfr}(t, f) df}$$
$$B^2(t) = \frac{\int_{-\infty}^{+\infty} f^2 \text{tfr}(t, f) df}{\int_{-\infty}^{+\infty} \text{tfr}(t, f) df} - f_m(t)^2;$$

for the *time moments*, and as

$$t_m(f) = \frac{\int_{-\infty}^{+\infty} t \text{tfr}(t, f) dt}{\int_{-\infty}^{+\infty} \text{tfr}(t, f) dt}$$

$$T^2(f) = \frac{\int_{-\infty}^{+\infty} t^2 \text{tfr}(t, f) dt}{\int_{-\infty}^{+\infty} \text{tfr}(t, f) dt} - t_m(f)^2;$$

for the *frequency moments*. They describe the averaged positions and spreads in time and in frequency of the signal. For some particular distributions, if the signal is considered in its analytic form, the first order moment in time also corresponds to the instantaneous frequency, and the first order moment in frequency to the group delay of the signal. These moments can be obtained numerically thanks to the functions `momttfr.m` and `momftfr.m`.

5.1.2 Marginals

It can also be interesting to consider the *marginal distributions* of a time-frequency representation. These marginals are defined as:

$$m_f(t) = \int_{-\infty}^{+\infty} \text{tfr}(t, f) df \quad \textit{time marginal}$$

$$m_t(f) = \int_{-\infty}^{+\infty} \text{tfr}(t, f) dt \quad \textit{frequency marginal}$$

and express, by integrating the representation along one variable, the repartition of the energy along the other variable. A natural constraint for a time-frequency distribution is that the time marginal corresponds to the instantaneous power of the signal, and that the frequency marginal corresponds to the energy spectral density:

$$m_f(t) = |x(t)|^2 \quad \text{and} \quad m_t(f) = |X(f)|^2.$$

The M-file `margtfr.m` computes the marginal distributions of a given time-frequency representation.

5.2 More on interferences: information on phase

The interference terms present in any quadratic time-frequency representation, even if they disturb the readability of the representation, contain some information about the analyzed signal. The precise knowledge of their

structure and construction rule is useful to interpret the information that they contain.

For instance, the interference terms contain some information about the phase of a signal. Let us consider the pseudo WVD of the superposition of two constant frequency modulations, with a phase shift between the two sinusoids. If we compare the pseudo WVD for different phase shifts, we can observe a time-sliding of the oscillating interferences. The M-file `movpwdph.m` illustrates this property (see fig. 5.1) :

```
>> M=movpwdph(128); movie(M,10);
```




Figure 5.1: Two simultaneous complex sinusoids analyzed by the pseudo-WVD: the position of the interferences depends on the phase-shift between the two components. These phase-shifts are respectively $\pi/4$, $3\pi/4$, $5\pi/4$ and $7\pi/4$

Each snapshot corresponds to the pseudo WVD with a different phase shift between the two components.

A second example of signature of the phase is given by the influence of a jump of phase in a signal analyzed by the (pseudo) Wigner-Ville distribution :

for instance, if we consider a constant frequency modulation presenting a jump of phase in its middle (see fig. 5.2) :

```
>> M=movpwjph(128,'C'); movie(M,10);
```




Figure 5.2: Complex sinusoid presenting a jump of phase in its middle, analyzed by the pseudo-WVD : the shape of the PWVD-pattern changes with the importance of the jump. These jumps of phase are respectively $\pi/4$, $\pi/2$, $3\pi/4$ and π

the pseudo WVD presents a pattern around the jump position which is all the more important since this jump of phase is close to π . This characteristic can be used to detect a jump of phase in a signal.

5.3 Renyi information

Another interesting information that one may need to know about an observed non-stationary signal is the number of elementary signals composing this observation. This also leads us to the following question: how much

separation between two elementary signals must one achieve in order to be able to conclude that there are two signals present rather than one ?

A solution to this problem is given by applying an information measure to a time-frequency distribution of the signal. Unfortunately, the well known Shannon information, defined as

$$I_x = - \int_{-\infty}^{+\infty} f(x) \log_2 f(x) dx$$

where $f(x)$ is the probability density function of x , can not be applied to some time-frequency distributions due to their negative values. The generalized form of information, which admits negative values in the distribution, will then be used. This information, known as *Renyi information*, is given by

$$R_x^\alpha = \frac{1}{1-\alpha} \log_2 \left\{ \int_{-\infty}^{+\infty} f^\alpha(x) dx \right\}$$

in the continuous case, where α is the order of the information. First order Renyi information ($\alpha = 1$) reduces to Shannon information. Third order Renyi information, applied to a time-frequency distribution $C_x(t, \nu)$, is defined as

$$R_C^3 = -\frac{1}{2} \log_2 \left\{ \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} C_x^3(t, \nu) dt d\nu \right\}.$$

The result produced by this measure is expressed in *bits*: if one elementary signal yields zero bit of information (2^0), then two well separated elementary signals will yield one bit of information (2^1), four well separated elementary signals will yield two bits of information (2^2), and so on. This can be observed by considering the WVD of one, two and then four elementary atoms, and then by applying the Renyi information on them. The file `renyi.m` computes this information measure :

```
>> sig=atoms(128,[64,0.25,20,1]);
>> [TFR,T,F]=tfrwv(sig);
>> R1=renyi(TFR,T,F)          -----> -0.2075

>> sig=atoms(128,[32,0.25,20,1;96,0.25,20,1]);
>> [TFR,T,F]=tfrwv(sig);
>> R2=renyi(TFR,T,F)          ----->  0.779

>> sig=atoms(128,[32,0.15,20,1;96,0.15,20,1;...
                 32,0.35,20,1;96,0.35,20,1]);
>> [TFR,T,F]=tfrwv(sig);
>> R3=renyi(TFR,T,F)          ----->  1.8029
```

We can see that if R is set to 0 for one elementary atom by subtracting R_1 , we obtain a result close to 1 for two atoms ($R_2 - R_1 = 0.99$) and close to 2 for four atoms ($R_3 - R_1 = 2.01$). If the components are less separated in the time-frequency plane, the information measure will be affected by the overlapping of the components or by the interference terms between them (see [WBI91] for more details on this analysis). In particular, it is possible to show that the Renyi information measure provides a good indication of the time separation at which the atoms are essentially resolved, with a better precision than with the time-bandwidth product.

5.4 Time-frequency analysis: help to decision

5.4.1 General considerations

The decision problem that one can have to solve when analyzing a signal is threefold :

- detect if an observed signal contains a given information (i.e. say, for a given false alarm probability, if *yes* or *no* the information is present) ;
- estimate the parameters of a signal that we know to be present in an observation ;
- classify a signal in one among different classes.

This problem, well known in theory in the general case, can be reconsidered when dealing with non-stationary signals, emphasized by the theory of time-frequency representations. Without going into details, it has been shown that some of the known optimal strategies of decision can be reformulated equivalently in the time-frequency plane (like the matched-filter with the WVD for example). This result is interesting for two reasons :

- on one hand, the time-frequency approach, compared to the classical one (formulated in the time-domain in general), usually provides a simpler interpretation of the decision test ;
- on the other hand, when the optimal solution for a given criterion is not known in the decision theory, the time-frequency analysis can be useful to formulate a sub-optimal solution based on the better comprehension of the analyzed signal (for example, a time-frequency detector can be easily modified to take into account variations of the non-stationary

signal to be detected, in order to improve the robustness of the detector).

The proposed solutions in the literature construct a decision test (statistic)

- either as a general time-frequency correlation between a time-frequency representation of the analyzed signal and some two dimensional template, constructed using the *a priori* information available on the signal,
- or by applying a transform on the TF representation of the analyzed signal, which brings to the fore some characteristic pattern of the signal to be detected (or estimated or classified), and by applying a test on this new space of decision. We consider in the following an example of such approach, for the problem of the detection and estimation of a linear frequency modulated signal embedded in some white gaussian noise.

5.4.2 An example : detection and estimation of linear FM signals

As we have seen in section 4.1.1, the WVD ideally concentrates the linear chirp signals in the time-frequency plane. Thus, the problem of detection and estimation of such a signal, which is not easily recognizable in the time-domain, is reduced to the problem of detection and estimation of a line in an image, which is a well known and easy-to-solve problem in pattern recognition. This can be done by using the Hough transform, dedicated to the detection of lines ([Bar95]).

The Hough transform for lines

Consider the polar parameterization of a line

$$x \cos \theta + y \sin \theta = \rho$$

(this parameterization is much more adapted to this problem than the Cartesian one). For each point (x, y) of an image I , the Hough transform associates a sinusoid in the plane (ρ, θ) , whose points have an amplitude equal to the intensity of the pixel (x, y) . So to all the points in I , the Hough transform associates a pencil of sinusoids which intersect themselves in the plane (ρ, θ) . In other words, the HT performs integrations along lines on the image I , and the value of each integral is affected to the point (ρ, θ) corresponding to the parameters of this line. Therefore, if on the image I some pixels with

high intensities are concentrated along a straight line, we will observe in the domain (ρ, θ) a peak whose coordinates are directly related to the parameters of the lines.

This method can be easily applied to other parametric curves, like hyperbola for example. This transform is computed in the file `ht1.m`.

The Wigner-Hough transform

When applying the Hough transform to the Wigner-Ville distribution of the signal

$$x(t) = e^{j2\pi(\nu_0 t + \beta/2t^2)} + n(t)$$

observed during an observation time T ($n(t)$ is a noise assumed white and gaussian), we obtain a new transform called the *Wigner-Hough transform* (WHT), whose expression is

$$\begin{aligned} WH_x(\nu_0, \beta) &= \int_T W_x(t, \nu_0 + \beta t) dt \\ &= \int_{-\infty}^{+\infty} \int_T x(t + \tau/2) x^*(t - \tau/2) e^{-j2\pi(\nu_0 + \beta t)\tau} dt d\tau \end{aligned} \quad (5.1)$$

The comparison of the WHT to a threshold is the proposed detection test, and the estimates of the unknown parameters ν_0 and β are given by the coordinates of the detected peak in the space of the parameters (ν_0, β) . Thanks to the unitarity property of the WVD (Moyal's formula), it is possible to show that this detection test is *asymptotically* the *optimal detector* (i.e. optimal when T tends to infinity). Besides, the *estimators* are *asymptotically efficient* (i.e. they asymptotically reach the Cramer-Rao lower bounds). Compared to the classical decision test usually used in this case, the generalized likelihood ratio test (GLRT), this method presents the following advantages in the case of multicomponent signals :

- it is free from the estimation of the initial phase and amplitude of each component, which usually do not bring any information, and
- its complexity do not increase with the number of components N_c , unlike the GLRT whose complexity increases linearly with N_c .

Here is an illustration of this decision test : first, we consider a linear chirp signal embedded in a white gaussian noise, with a 1 dB signal-to-noise ratio :

```
>> N=64; sig=sigmerge(fmlin(N,0,0.3),noisecg(N),1);
```

Now, if we analyze it with the WVD followed by the Hough transform (see fig. 5.3 and 5.4),

```
>> tfr=tfrwv(sig); contour(tfr,5); grid
>> ht1(tfr,N,N,1);
```




Figure 5.3: WVD of a noisy chirp signal (SNR=1 dB) : while the chirp is hardly readable in the time-representation, the line still clearly appear in the WVD

we obtain, in the parameters' space (ρ, θ) , a peak representing the chirp signal, significantly more energetic than the other peaks corresponding to the noise. The decision test is then very simple : it consists in applying a threshold on this representation, positioned according to a detection criterion ; if the peak is higher than the threshold, then the chirp is said to be present, and the coordinates of that peak $(\hat{\rho}, \hat{\theta})$ provide estimates of the chirp parameters (the change from $(\hat{\rho}, \hat{\theta})$ to $(\hat{\nu}_0, \hat{\beta})$ corresponds to the change from polar to Cartesian coordinates).

In the case of a multi-component signal, the problem of interference terms appear. However, due to the oscillating structure of these terms, the integration (5.1) operated by the Hough transform on the WVD will attenuate




Figure 5.4: Wigner-Hough transform of the previous noisy chirp: the peak corresponds to the chirp signal (and the side-lobes to the noise), and its coordinates give estimators of the chirp parameters. The detection test consists in comparing this peak to a threshold (threshold fixed by the chosen criterion)

them. This can be observed on the following example: we superpose two chirp signals with different initial frequencies and sweep rates (see fig. 5.5 and 5.6):

```
>> sig=sigmerge(fmlin(N,0,0.4),fmlin(N,0.3,0.5),1);
>> tfr=tfrwv(sig); contour(tfr,5); grid
>> htl(tfr,N,N,1);
```

We can see that the components are well separated in the parameter space, in spite of the use of a nonlinearity in the WHT. Again, the coordinates of the two peaks provide estimates of the different parameters.

5.5 Analysis of local singularities

If the time-frequency representations are useful to bring to the fore the progression with time of the frequency of a signal, the time-scale representations are more adapted to the analysis of irregular structures and singularities, or of signals presenting self-similarities (such as fractional Brownian




Figure 5.5: WVD of two simultaneous chirp signals : interference terms appear between the two components

motion, [GF92]). We give in the following such an example with the analysis of *local singularities*, thanks to the scalogram and the Unterberger distribution.

The local regularity of a signal can be characterized by its *Holder* (or *Lipschitz* or *scaling*) *exponent*: for a signal $x(t)$ which is *uniformly Holder* H , there exists a constant C such that

$$|x(s) - x(t)| \leq C |s - t|^H, \quad 0 < H < 1.$$

H then represents the *exponent of regularity* of the signal. If we consider the wavelet transform $T_x(t, a; \Psi)$ of this signal, with an analyzing wavelet Ψ such that $t \Psi(t)$ is absolutely integrable, then one can show that

$$\begin{aligned} |T_x(t, a; \Psi)| &\leq C |a|^{H+1/2} \int_{-\infty}^{+\infty} |t|^H |\Psi(t)| dt \\ &= O(|a|^{H+1/2}) \quad \forall t, \end{aligned}$$

or, in terms of scalogram and behavior when a tends to 0,

$$E \left[|T_x(t, a; \Psi)|^2 \right] \sim |a|^{2H+1}, \quad a \rightarrow 0.$$




Figure 5.6: Wigner-Hough transform of the two-component chirp signal : two main peaks are present, characterizing the two chirp components, while the cross terms present in the WVD only introduce small side-lobes in the Wigner-Hough transform

where $E[.]$ refers to the expectation. This means that the regularity of the signal can be recovered from the behavior of its scalogram at small scales, and it is possible to show that the reciprocal is true.

Since they are time-dependent in nature, the wavelet-based techniques also allow an estimation of the local regularity of a signal. In some sense, time-scale methods offer in this respect a framework similar to the one provided by time-frequency analysis for tracking the time evolution of spectral features. Indeed, if we now have, at a given time t_0 ,

$$|x(t_0 + \tau) - x(t_0)| \leq C |\tau|^{H(t_0)}, \quad 0 < H(t_0) < 1, \quad (5.2)$$

then we can establish the inequality

$$\begin{aligned} |T_x(t, a; \Psi)| &\leq C |a|^{H(t_0)+1/2} \int_{-\infty}^{+\infty} |t|^{H(t_0)} |\Psi(t)| dt \\ &\quad + C |t - t_0|^{H(t_0)} \int_{-\infty}^{+\infty} |\Psi(t)| dt \\ &= O(|a|^{H(t_0)+1/2} + |t - t_0|^{H(t_0)}). \end{aligned}$$

We then obtain an image of the signal's regularity at the small scales of its wavelet transform (or scalogram), but accompanied with a time localization.

The reciprocal is also true, which means that an appropriate decrease of the wavelet (scalogram) coefficients in a cone-shaped region of the time-frequency plane allows one to estimate the local regularity of a signal.

If we further impose to condition (5.2) that the signal presents an asymptotic spectral decrease,

$$X(\nu) \sim |\nu|^{-(1+2H(t_0))} e^{j2\pi\nu t_0} \quad \text{for } |\nu| \rightarrow \infty,$$

then we have the following approximation for the active Unterberger distribution :

$$U_x(t, a) \sim |a|^{2(1+H(t_0))} \delta(t - t_0), \quad a \rightarrow 0.$$

Thus, the Unterberger distribution follows a law along scales which gives access to the strength of the singularity (H), and along time to the localization of this singularity.

The file `holder.m` estimates the Holder exponent of any signal from an affine time-frequency representation of it.

o *Example*

For instance, we consider a 64-points Lipschitz singularity (see `anasing.m`) of strength $H = 0$, centered at $t_0 = 32$,

```
>> sig=anasing(64);
```

and we analyze it with the scalogram (Morlet wavelet with half-length = 4, see fig. 5.7),

```
>> [tfr,t,f]=tfrscalogram(sig,1:64,4,0.01,0.5,256,1);
```

The time-localization of the singularity can be clearly estimated from the scalogram distribution at small scales :

```
>> H=holber(tfr,f,1,256,32) -----> H=-0.0381
```

If we now consider a singularity of strength $H=-0.5$ (see fig. 5.8),

```
>> sig=anasing(64,32,-0.5);
>> [tfr,t,f]=tfrscalogram(sig,1:64,4,0.01,0.5,256,1);
```

we notice the different behavior of the scalogram along scales, whose decrease is characteristic of the strength H . The estimation of the Holder exponent at $t = 32$ gives :

```
>> H=holber(tfr,f,1,256,32) -----> H=-0.5107
```

which is close to 0.5.

The same conclusions can be observed from the active Unterberger distribution.




Figure 5.7: Scalogram of a Lipschitz singularity at time $t = 32$, of strength $H = 0$




Figure 5.8: Scalogram of a Lipschitz singularity at time $t = 32$, of strength $H = -0.5$

GNU Free Documentation License

Version 1.2, November 2002

Copyright ©2000,2001,2002 Free Software Foundation, Inc.

51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The **”Document”**, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as **”you”**. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A **”Modified Version”** of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A **”Secondary Section”** is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The **”Invariant Sections”** are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The **”Cover Texts”** are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A **”Transparent”** copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not **”Transparent”** is called **”Opaque”**.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "**Title Page**" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

A section "**Entitled XYZ**" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "**Acknowledgements**", "**Dedications**", "**Endorsements**", or "**History**".) To "**Preserve the Title**" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if

there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.

- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements".

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or

any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright ©YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled “GNU Free Documentation License”.

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the ”with...Texts.” line with this:

with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST.

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

Bibliography

- [AF94] F. Auger and P. Flandrin. The why and how of time-frequency reassignment. *IEEE International Symposium on Time-Frequency and Time-Scale Analysis*, pages 197–200, 1994. Philadelphia.
- [AF95] F. Auger and P. Flandrin. Improving the Readability of Time-Frequency and Time-Scale Representations by the Reassignment Method. *IEEE Transactions on Signal Processing*, 43(5):1068–89, 1995.
- [AGT91] L. Auslander, I. Gertner, and R. Tolimieri. The discrete zak transform application to time-frequency analysis and synthesis of nonstationary signals. *IEEE Transactions on Signal Processing*, 39(4):825–835, 1991.
- [Aug91] F. Auger. *Representations temps-frequence des signaux non-stationnaires: synthese et contributions*. PhD thesis, Ecole Centrale de Nantes, France, 1991.
- [Bar95] S. Barbarossa. Analysis of Multicomponent LFM Signals by a Combined Wigner-Hough Transform. *IEEE Transactions on Signal Processing*, 43(6), June 1995.
- [BB92] J. Bertrand and P. Bertrand. A class of affine wigner functions with extended covariance properties. *J. Math. Phys.*, 33(7), 1992.
- [Coh89] L. Cohen. Time-Frequency Distributions - A Review. *Proceedings of the IEEE*, 77(7):941–980, 1989.
- [Dau92] I. Daubechies. *Ten Lectures on Wavelets*. SIAM, 1992.
- [Fla86] P. Flandrin. On Detection-Estimation Procedures in the Time-Frequency Plane. In *International Conference on Acoustics, Speech and Signal Processing*, pages 43.5.1–4, 1986.

- [Fla93] P. Flandrin. *Temps-fréquence*. Hermès, 1993. Trait des Nouvelles Technologies, srie Traitement du Signal.
- [GB96] P. Gonalvs and R. Baraniuk. Pseudo affine wigner distributions and kernel formulation. *Submitted to IEEE Transactions on Signal Processing*, 1996.
- [GF92] P. Gonçalvès and P. Flandrin. Scaling exponents estimation from time-scale energy distributions. In *IEEE Int. Conf. on Acoust., Speech and Signal Proc. ICASSP-92*, pages V.157–V.160, San Francisco (CA), 1992.
- [Gon93] P. Gonalvs. *Représentations temps-fréquence et temps-échelle bilinéaires: synthèse et contributions*. PhD thesis, Institut Polytechnique de Grenoble - Laboratoire de Physique de l'École Normale Supérieure de Lyon, France, 1993.
- [Har78] F. Harris. On the use of windows for harmonic analysis with the discrete fourier transform. In *Proceedings IEEE*, volume 66, pages 51–83, 1978.
- [HBB92] F. Hlawatsch and F. Boudreaux-Bartels. Linear and Quadratic Time-Frequency Signal Representations. *IEEE SP Magazine*, pages 21–67, 1992.
- [Hla91] F. Hlawatsch. Time-frequency methods for signal processing. In *Technical Report 1291-0001, Dept of Electrical Engineering, University of Rhode Island*, 1991.
- [KdVG76] K. Kodéra, C. de Villedary, and R. Gendrin. A New Method for the Numerical Analysis of Time-Varying Signals with Small BT Values. *Phys. Earth Planet. Interiors*, 12:142–150, 1976.
- [Ova94] J-P. Ovarlez. La transformation de mellin et l'analyse des signaux large-bande. *Colloque Temps-Fréquence, Ondelettes et Multirésolution, INSA-Lyon*, pages 13.1–13.8, 9-11 Mars 1994.
- [RF92] O. Rioul and P. Flandrin. Time-scale distributions : A general class extending wavelet transform. *TRANSSIG*, 40(7):1746–57, July 1992.
- [WBI91] W. Williams, M. Brown, and A. Hero III. Uncertainty, information, time-frequency distributions. *SPIE Advanced Signal Processing Algorithms, Architectures and Implementations II*, 1566:144–156, 1991.