

A Quantitative Measure of Mono-Componentness for Time-Frequency Analysis

A Dissertation Presented for the
Doctor of Philosophy
Degree
The University of Tennessee, Knoxville

Austin P. Albright

December 2016

© by Austin P. Albright, 2016
All Rights Reserved.

Dedication

To the glory of my Lord and Savior, Jesus Christ.

To my beautiful, patient wife Melissa,

and to our four children: William, Molly, Maggie, and Weston.

Acknowledgments

First and most importantly, I am deeply indebted to my wife Melissa for her support, patience and encouragement. It wasn't easy and it wasn't fun but she did make it. She deserves a PhD as much as I do for everything it takes to deal with four small children and a husband going back to school after seven years. I am also thankful for my children Will, Molly, Maggie, and Weston. There is nothing like a herd of small children to get your mind off your research problems. Their little faces greeting me when I get home is the perfect antidote to beating my head against the PhD wall all day, every day. Where I am today is in no small part due to their love and support... especially Melissas.

I also want to thank Dr. Hairong Qi for her phenomenal and unwavering support throughout this endeavor. Without her support I never would have been able to return after all those years between completing my masters and finishing this doctorate. When Dr. Qi accepts you as one of her graduate students, she real cares about you, your success, and even your wife and kids in my case. She played no small part in this achievement. Dr. Mark Buckner also went above and beyond the call of committee member by meeting with me once a week to help keep me on track and from following rabbit trails or getting tangled up in the trivial... for more than a week at a time. Those meetings were especially critical during the summer months. Without Dr. Qi and Dr. Buckner I never would have finished this dissertation in the relatively small window of two years.

There are two other people that played a critical role in this process that I feel

are often overlooked despite the roles they play. They are Ms. Dana Bryson in the Electrical Engineering and Computer Science (EECS) department and Ms. Sylvia Miller in the UTK Graduate School office. Due to all the changes in the EECS program and the length of time between when I started accumulating credits towards Ph.D. and returned to finish, a lot changed. In fact, if it were not for Dana and Sylvia this entire endeavor would have failed. It was their guidance and support in navigating the treacherous and critical paperwork in the approval process for returning after so many years away, which allowed me to even dream of returning to finish my doctorate. Without those approvals this would have all been for nothing and died on the vine.

I would like to thank the members of my committee: Dr. Mark Buckner, Dr. Bella R. Upadhyaya, and Dr. Husheng Li. I greatly appreciate their time and patience. I want to thank all of the AICIP crew for their friendship and company during this process. Likewise, I want to thank the Cognitive Radio “tribe” at Oak Ridge National Laboratory (ORNL). They are my co-workers and closest friends before, during, and after this PhD/sabbatical adventure. Not the least of whom was my group leader Paul Ewing. A great man who died too soon.

Also, I need to acknowledge ORNL which allowed me to take an educational sabbatical to focus full-time on the research, development, and writing of this dissertation.

Finally, because this is my acknowledgment section and I can say whatever I want... This is for my four little buddies, who always want a “funny story” made up for them at bedtime. All of which are required, by them, to end with “The Weasel Dance” the accompanying lyrics of which are... Weasel, weasel, weasel! Weasel, weasel, weasel! Mink! Ferret!

Abstract

Joint time-frequency (TF) analysis is an ideal method for analyzing non-stationary signals, but is challenging to use leading to it often being neglected. The exceptions being the short-time Fourier transform (STFT) and spectrogram. Even then, the inability to have simultaneously high time and frequency resolution is a frustrating issue with the STFT and spectrogram. However, there is a family of joint TF analysis techniques that do have simultaneously high time and frequency resolution – the quadratic TF distribution (QTFD) family. Unfortunately, QTFDs are often more troublesome than beneficial. The issue is interference/cross-terms that causes these methods to become so difficult to use. They require that the “proper” joint distribution be selected based on information that is typically unavailable for real-world signals. However, QTFDs do not produce cross-terms when applied to a mono-component signal.

Clearly, determining the mono-componentness of a signal provides a key piece of information. However, until now, the means for determining if a signal is a mono-component or a multi-component has been to choose a QTFD, generate the TF representation (TFR), and visually examine it. The work presented here provides a method for quantitatively determining if a signal is a mono-component. This new capability provides an important step towards finally allowing QTFDs to be used on multi-component signals, while producing few to no interference terms through enabling the use of the quadratic superposition property. The focus of this work is on establishing the legitimacy for “measuring” mono-componentness along with its

algorithmic implementation. Several applications are presented, such as quantifying the quality of the decomposition results produced by the blind decomposition algorithm, Empirical Mode Decomposition (EMD).

The mono-componentness measure not only provides an objective means to validate the outcome of a decomposition algorithm, it also provides a practical, quantitative metric for their comparison. More importantly, this quantitative measurement encapsulates mono-componentness in a form which can actually be incorporated in the design of decomposition algorithms as a viable condition/constraint so that true mono-components could be extracted. Incorporating the mono-component measure into a decomposition algorithm will eventually allow interference free TFRs to be calculated from multi-component signals without requiring prior knowledge.

Table of Contents

1	Introduction	1
1.1	Time-Frequency Analysis	1
1.2	The Motivating Challenge	5
1.3	Inspiration Strikes	7
1.4	The Contributions	9
1.5	Organization	11
2	Background and Prior Work	13
2.1	Terminology	13
2.2	A Brief Historical Overview	14
2.3	Multi-component and Mono-component Signals	17
2.4	Ridges	20
2.5	The Wigner-Ville Distribution	23
2.5.1	Discrete Wigner-Ville Distribution	25
2.6	Interference Terms in the Wigner-Ville Distribution	26
2.7	Summary	34
3	The Measurement of Mono-Componentness	35
3.1	Wigner-Ville distribution of a Mono-component	37
3.2	The Algorithm	40
3.3	Ridge Detection	42
3.3.1	Crazy Climbers Algorithm	44

3.3.2	Chaining	45
3.3.3	Gap Filling	47
3.4	Implementation Considerations	48
3.4.1	Auto-WVD or Cross-WVD?	48
3.4.2	Sensitivity and False Positives	48
3.4.3	Handling Large Numbers of Ridge Detections	53
3.5	Limitations and Usage	54
3.6	Mono-component Measurement Case Studies	55
3.6.1	True Mono-component Case	56
3.6.2	Multi-component Case	58
4	Applications and Analysis using the Mono-componentness Measure	61
4.1	Single-Channel Decompositions	62
4.1.1	Empirical Mode Decomposition	63
4.1.2	Known Limitations of EMD	66
4.2	Empirical Multi-Component Decomposition	75
4.2.1	An Advantageous Combination	76
4.2.2	An Analysis of the Mono-componentness of EMD's IMFs	77
4.2.3	Problematic Cases	90
4.3	Quantitative Measure of the Quality for the Hilbert-Huang Transform Spectrum	90
5	Discussion	93
5.1	In Summary	93
5.2	Future Work	95
5.2.1	Algorithmic Inner Interference Detection	96
5.2.2	Improved or New Decomposition Method	97
5.2.3	Computational Complexity	99
5.2.4	Parameter Values	99
5.2.5	A Better Ridge Detector	100

Bibliography	102
Appendix A Software Resources	115
Vita	116

List of Tables

3.1	The starting IF, f_0 , and ending IF, f_{end} , values used to generate each of the three linear FM mono-components.	58
-----	---	----

List of Figures

1.1	Illustration of shortcomings of time or frequency only analysis	3
1.2	Illustration of advantage of joint time-frequency analysis	4
2.1	In these idealized TF plots each component of a signal is perfectly concentrated to form a single ridge.	20
2.2	The “foothills scenario” can be seen in this zoomed in view of the TFR generated via the WVD for a simple mono-component signal.	21
2.3	Edge detection versus ridge detection in a TFR	22
2.4	Outer interference in the WVD	30
2.5	Inner interference in the WVD	32
3.1	Flowchart of mono-componentness measure.	41
3.2	The TFR, Canny edge detection, and ridge detection results of a 2-component multi-component signal	43
3.3	TFR and Crazy Climbers occupancy grid example	46
3.4	Illustrating the detectability of a cross-term even when one of the auto-terms is not	51
3.5	A real-world illustration of how the detectability of the cross-terms, due to their higher amplitude, provides information on the existence of an all but imperceptible auto-term.	52
3.6	Intermediary results of applying the mono-componentness measure a true mono-component	57

3.7	Intermediary results of applying the mono-componentness measure to a multi-component signal	59
4.1	Example of mode-mixing in EMD results	67
4.2	Two-tone separability score for a EMD using a 10 fixed iterations . .	70
4.3	Two-tone separability score for a EMD using a 100 fixed iterations . .	71
4.4	The two-tone separability score for EMD when using a fixed 10-iterations per IMF sifting viewed in 3-D.	72
4.5	The two-tone separability score for EMD when using a fixed 100-iterations per IMF sifting viewed in 3-D.	72
4.6	The two-tone separability score for EMD when using the “classic” tolerance based stopping criteria, $[\theta_1 = 0.05, \theta_2 = 0.5, \alpha = 0.05]$	73
4.7	The two-tone separability score for EMD when using the “classic” tolerance based stopping criteria viewed in 3-D.	73
4.8	The time domain plot of the 3-component test signal.	79
4.9	The time domain plot of the input to EMD (topmost pane), followed in descending order by the first IMF extracted ending with the residue (bottommost pane).	80
4.10	WVD of the 3-component test signal and the individual mono-components that compose it	81
4.11	Intermediary results from applying the mono-component measure to the decomposition results from EMD	82
4.12	IMFs extracted when using EMD with 10 fixed iterations per IMF . .	84
4.13	Intermediary results from applying the mono-component measure to the decomposition results from EMD using 10-iterations	85
4.14	IMFs extracted when using EMD with 100 fixed iterations per IMF .	86
4.15	Intermediary results from applying the mono-component measure to the decomposition results from EMD using 100-iterations	87

4.16	IMFs extracted when using EMCD with a maximum of a 100 iterations per IMF	88
4.17	Intermediary results from applying the mono-component measure to the decomposition results from EMCD using 100-iterations and the mono-component measure	89
4.18	The HHT from EMD and the HHT from the ground truth	92

Nomenclature

CC	Crazy Climber
CWT	Continuous Wavelet Transform
DWVD	Discrete Wigner-Ville Distribution
EMCD	Empirical Multi-Component Decomposition
EMD	Empirical Mode Decomposition
FFT	Fast Fourier Transform
IA	Instantaneous Amplitude
IAF	Instantaneous Auto-correlation Function
ICA	Independent Component Analysis
IF	Instantaneous Frequency
IMF	Intrinsic Mode Function
LAF	Local Auto-correlation Function
OFDM	Orthogonal Frequency-Division Multiplexing
QTFD	Quadratic Time-Frequency Distribution
STFT	Short-Time Fourier Transform

TF	Time-Frequency
TFR	Time-Frequency Representation
WVD	Wigner-Ville Distribution

Chapter 1

Introduction

1.1 Time-Frequency Analysis

Time-frequency analysis is an intuitive everyday part of our lives. We want to know what happened and when. We make plans based on when and where. At work we want to know what to do and when to do it. The same desires and needs apply to how we examine the world around us. Whether it is historical, political, financial, environmental, or electrical, knowing both the what and the when of an “event” is highly useful and informative. In this document, the focus is specifically on time-frequency analysis in the field of signal processing. In this field, time is the “when” and frequency is the “what.”

Why joint time-frequency (TF) analysis? What is wrong with just spectral analysis (Fourier) and time-series analysis? It is not that there is anything wrong with time domain or frequency domain analysis. They are the fundamental building blocks of any type of analysis after all. But there is so much more that can be learned if we work in the joint time-frequency plane. The following simple example makes this profoundly clear.

Using just two simple signals it is easy to see the shortcomings of using disjointed analysis, in regards to the information yielded. Again, it is not that the two disjointed

domains are bad and should not be used. It is that by themselves information is missing.


For clarity, let's name one signal α and the other β . The α signal is simply half a second of null (zero), followed by three seconds of a 6 Hz sinusoid, then three seconds of a 10 Hz sinusoid, and ending with another a half second of null. The β signal is exactly the same as α except that the three seconds of a 10 Hz sinusoid occur first followed by three seconds of a 6 Hz sinusoid. Equation (1.1) and Eq. (1.2) show the formulas for α and β , respectively.

$$\alpha = \begin{cases} 0, & 0.0 \leq t \leq 0.5, \\ \cos(6 \cdot 2\pi t), & 0.5 < t \leq 3.5 \\ \cos(10 \cdot 2\pi t), & 3.5 < t \leq 6.5 \\ 0, & 6.5 < t \leq 7.0, \end{cases} \quad (1.1)$$


$$\beta = \begin{cases} 0, & 0.0 \leq t \leq 0.5, \\ \cos(10 \cdot 2\pi t), & 0.5 < t \leq 3.5 \\ \cos(6 \cdot 2\pi t), & 3.5 < t \leq 6.5 \\ 0, & 6.5 < t \leq 7.0, \end{cases} \quad (1.2)$$

Looking at the time domain plots, Figure 1.1(a,c), we can visually tell that there is a difference between the two signals and when the change occurs, but not what the change was. In Figure 1.1(b,d), we see the frequency domain representation of the two signals which shows what was in the signals, but not when it occurred. In fact, the two signals are all but indistinguishable in the frequency domain. In contrast, by using a joint TF method to analyze the α , Figure 1.2(a), and β signals, Figure 1.2(b), we see the power and advantages of performing an analysis in the joint domain immediately.


The advantages of using time-frequency analysis methods makes intuitive sense.




(a) The α signal in the time domain.



(b) Power spectral density of the α signal.




(c) The β signal in the time domain.



(d) Power spectral density of the β signal.

Figure 1.1: These figures illustrate the major short coming of analyzing signals in a single domain. In the time domain plots of (a) and (c) it is clear that there is a change in each signal, but what those changes are is not available in the time domain. In plots (b) and (d), the frequency information is available, but there is no timing context at all, i.e., start time, stop time, duration.



(a) A joint time-frequency representation of the α signal. (b) A joint time-frequency representation of the β signal.

Figure 1.2: These two figures illustrate the advantage of joint time-frequency analysis. It is immediately clear what frequencies occurred, when they occurred, and for how long they persisted.

Where do we start if all that is known about the signal is that an event was captured – how do we begin to analyze it? We need the same basic information used to go about our daily lives. When and where. When did the “event” occur? Where in the spectrum was it? TF analysis has the ability to answers both of these questions.

There are two general categories of TF techniques: linear, and quadratic. The linear methods are the fairly straightforward to use but suffer one critical shortcoming – their resolutions are limited. The most well known of the linear methods is the short time Fourier transform (STFT). It is not possible to have simultaneously high resolution in both time and frequency using a linear method. Consider for a moment if we planned meetings using the same “resolution” as provided by STFT. It would be like organizing a meeting where the time of the meeting is provided down to the minute, but location is no more specific than a city’s name or providing the location down to the floor and room but only giving the month when the meeting is to occur. These analogies, though a bit extreme, are exactly the resolution issue faced when using linear methods.

Naturally, a TF analysis which yields high resolution in time and frequency

simultaneously is the ideal solution. This desire can be fulfilled through the use of a non-linear time-frequency method. Unfortunately, the old adage “There is no such thing as a free lunch” is still as true as ever. The quadratic methods do provide superior resolution in both time and frequency simultaneously, but are more challenging to use.

1.2 The Motivating Challenge

Much of the broadband, high resolution data that is regularly collected these days are prime candidates for more advanced analysis methods than simply Fourier. Especially, with the additional consideration that much of the real-world is populated by non-stationary signals. Fast, burst-like transient signals are rarely going to be stationary and definitely not when the scope is the entire duration of the recording. Finding such signals in captures made at high sampling rates and using the wide acquisition bandwidths available today are just asking to be analyzed in the joint TF domain. Not only does analysis in the TF domain allow localization in time and frequency it is fully capable of operating on non-stationary signals. On the other hand, we are frequently confronted with multi-component signal analysis problems involving signals that little is known about, and where assuming it is stationary and periodic is a risky proposition. Whether the signal’s source is a recording of a digital bus, an over-the-air transmission, environmental sampling, speech, or even an exotic physics experiment — multi-component signals are the norm, not the exception.

The problem facing the practical, real-world, useful application of quadratic TF analysis techniques essentially boils down to one issue — interference terms between the multiple components of a signal. The results produced using quadratic formulations suffer from the presence of interference terms whenever a multi-component signal is analyzed (e.g., speech, ultrasound, biological signals, seismic, etc.). While the time and frequency resolutions of quadratic methods are greatly improved over linear methods, such as the STFT and the continuous wavelet

transform (CWT), the real-world usefulness of the quadratic methods are greatly hindered by this cross-term problem [1, 2, 3, 4, 5, 6].

Approaches to reduce or eliminate the interference terms in quadratic TF methods have been developed that make use of various kernels. This in itself creates an equally challenging issue, that in order to select the best TF method for a multi-component signal, a significant level of information about the signal is needed to make an informed selection. Information that is typically one of the *desired* outcomes of the TF analysis which is not available via time-series or Fourier analysis for non-stationary, multi-component signals, such as the instantaneous frequency (IF) law. That the selection of the most appropriate TF analysis method is dependent on the information we hope to obtain from the analysis has a profoundly negative impact on the usability of advanced TF methods on multi-component signals – and of course don't forget the abundant presence of cross-terms.

In effect, a circular dependency exists. Our knowledge of the multi-component signal's contents depends on having high resolution in both time and frequency, which necessitates using a quadratic TF method. Because it is a multi-component signal there will be cross-terms. The cross-terms need to be reduced, or eliminated if possible, which requires that the most appropriate quadratic TF method should be selected. To make the most appropriate selection, knowledge of the signal is needed that would be provided via a quadratic TF method... and around and around it goes.

In practice, this sort of circular problem results in using a number of different quadratic TF methods to analyze an unknown multi-component signal; where a human in the loop must then compare the results and parameter variations attempted, i.e., kernel parameters, window durations, etc., and make a selection as to which is the best. This in itself is risky since what is a cross-term, with no real-world physical meaning, and what is an auto-term is by no means readily distinguishable. It can also be quite time consuming to calculate what could be dozens, even hundreds, of TF results. Especially, if the signal was recorded using a high sampling rate*.

*By significant duration, the author is referring to situations such as sampling at several hundred

Clearly, we need to pick the best suited quadratic method for the analysis so as to reduce or eliminate the cross-terms as much as possible, while simultaneously attempting to maintain as high a time and frequency resolution as possible – two objectives that oppose each other. To eliminate the cross-terms we need to apply smoothing. Whether the smoothing occurs through filtering or the convolution of a kernel, the end effect is the same – a reduction in either the time or frequency resolution or both. So the resolution issue is really a cross-term issue. If there were not cross-terms we would not need to perform additional operations, such as smoothing, that result in reduced resolution. Interestingly enough, it turns out that the cross-term issues can be pinpointed to a single source — the multi-component signal.

There is another possibility for dealing with the cross-terms from multi-component signals. It has the advantage of not requiring that the “perfect” method be somehow selected, but it does require that *all* the fundamental components that make up the multi-component signal being analyzed are known. If such omniscient knowledge is available then it becomes possible to use a quadratic TF method while eliminating, in most cases, all of the cross-terms.

1.3 Inspiration Strikes

After extensively searching, reviewing, and investigating the vast amounts of material that exist on quadratic time-frequency methods, their cross-term problem, and the associated resolution trade-off it became apparent that the cross-term problem is a mathematical Mt. Everest. On one side is the multi-component signal we have. On the other side is the goal of a simultaneous high resolution in both domains without interference terms and sitting in the middle is the cross-term Mount Everest. People have been attempting to move the mountain in order to make a path between what we have (the multi-component signal) and what we want (high resolution without MHz all the way to multi-GHz for even a few seconds).

cross-terms) for decades.

It was through all this investigation that the spark of inspiration ignited. By recognizing the monumental task of mountain moving, even with a thorough mapping of the terrain already available thanks to the decades of work others have contributed, a different perspective emerged.

Do we have to move the mountain? What if we could obtain a set of mono-components and analyze them independently instead of analyzing the multi-component as a whole? Could we just walk around the cross-term mountain? The route may be a bit longer than if the mountain was moved, but the mountain has not been moved yet and appears to be staying for the foreseeable future. In the end, walking around the cross-term mountain still gets you to the goal.

By assessing the available knowledge and avoiding the known uncertainties the spark quickly grew into a workable approach. By taking advantage of the capabilities developed in other fields a path was developed that leverages the prolific and powerful improvements in computational resources available today, along with modern decomposition algorithms, and the thorough understanding of what produces cross-terms .

Eventually, informed by the knowledge that for quadratic TF methods cross-terms come from multi-component signals, but not from mono-components[†], a unique strategy of seeking the mono-components of a multi-component signal was formulated. We can effectively lower the level of uncertainty with respect to choosing a quadratic method and circumnavigate the cross-term mountain using modern computational resources and a unique perspective of the problem.

Instead of the traditional focus on the multi-component signal and the cross-terms resulting from applying a quadratic TF method to it; we focus instead on how to acquire a set of *valid* mono-components which when summed equal the original multi-component signal. The reason being so that a single quadratic method (which

[†]Inner interference, Section 2.6 is treated as a special case in this document, but in short is a situation where cross-terms are present but the signal is a true mono-component.

possesses the quadratic superposition property) could be used consistently on each mono-component rather than facing the struggle of picking the “best” method and attempting to generate a high resolution TF result with minimal blurring and cross-terms. This new perspective naturally leads to the questions:

- What is a mono-component?
- Can mono-componentness be identified?
 - Algorithmically?
 - Objectively?
 - Repeatably?
 - Without prior knowledge?

Ultimately, the desire to avoid the problematic selection of a quadratic method exposed the more fundamental need for determining if a signal is truly a mono-component or not. Determining if a signal is mono-component or not is an unaddressed need and challenge. It is this challenge to which a solution has been developed and will be presented.

1.4 The Contributions

The significant contribution to the field of time-frequency signal processing presented in this work is not in the form of a new TF method or design principles, but instead a new fundamental tool that provides a means for an important attribute of a signal to be determined. This allows you to know you have a multi-component signal (something typically safely assumed, but still it is by assumption not fact), to know if a single is a true mono-component, and it allows you to know if the decomposition results of a multi-component signal produced all mono-components, some mono-components, or no mono-components.

Knowing where you are starting from is a valuable decision making tool. Knowing if you are working with a set of mono-components creates a valuable new starting point. Suddenly, the prospects of using superposition to generate a high resolution TF representation without the trials and risk of selecting a kernel and any required kernel parameter tuning is within reach. Not only that, it provides a valuable metric which can be used as a criteria in the design of decomposition algorithms, an objective function, an error metric, as a post-processing test of the results, and so on.

There is currently *no* means of any kind for determining if a signal is a mono-component other than by generating a TF representation of the signal followed by visual examination. My method is the state-of-the-art as it is the only of its kind.

My unique contributions as a result of this work are:

- An objective measure for determining if a signal is a mono-component.

The additional contributions enabled solely due to the ability to algorithmically determine if a signal is a mono-component are:

- Aid the design and/or improvement of blind single-channel decomposition algorithms.
 - They can be designed around producing valid mono-component signals.
 - The results can be quantified and the thoroughness of the decomposition results characterized.
- The potential to produce a cross-term free Wigner-Ville distribution through superposition.
 - A critical piece of a practical solution is the knowledge of whether a “potential” mono-component is actually a mono-component.

A means of determining mono-componentness holds additional benefits when integrated into single-channel decomposition algorithms. For there is an issue with single-channel decomposition algorithms, which is that while they provide a

set of components for a given input, they make no guarantees that the resulting signals are mono-components. What has been missing from all of the single-channel decomposition methods available at this time, e.g., Empirical Mode Decomposition (EMD), is a way to objectively and reliably determine if a resulting “potential mono-component” is actually a mono-component. Yet another contribution of this work is:

- A new EMD-based blind single-channel decomposition was created.
 - By combining the measure with EMD, any intrinsic mode function (IMF) that passes the measure will be a true mono-component.

An algorithmic mono-componentness measure opens the door to the safe application of machine learning techniques for task such as the reliable extraction of TF features. The ability to process long time-series and evaluate them without having to display them as images for a human analyst; a task steeped in issues that is even further confounded by the issues of visualizing the results. The work presented here focuses specifically on this first of its kind, objective, algorithmic method for determining the mono-componentness of a signal and how such a “measure” is legitimate.

1.5 Organization

The interrelated nature of the cross-term problem and its source makes it difficult to present them as fully independent pieces of material. Despite this, it is easier to try to maintain a step-by-step logical train-of-thought by presenting the material as independent sections, but as a forewarning there will be some “bleed-over” of material between sections.

Chapter 2 begins by addressing a few terminology usages that are unique to the TF field, followed by a brief history of the Wigner-Ville distribution. It then covers the interrelated elements, multi-component signals, mono-component signals, the

Wigner-Ville distribution, the resulting interference issue that occurs when the WVD is calculated of a multi-component signal (or using any quadratic TF distribution for that matter). Along the way, what exactly is meant by ridge detection is covered as this provides the context for much of the work in Chapter 3.

In Chapter 3, the primary novel contribution of this work is presented focusing specifically on not only how mono-componentness can be quantified but how it can be done so elegantly and legitimately. Before transitioning to several examples of the new found power of an objective, algorithmic means for determining if a signal is a mono-component, two examples of its use and its outcomes are presented.

In Chapter 4, we go from getting to know the mono-component measure to putting it to work for us by showing how it can be used to validate the results from a single channel blind decomposition algorithm, to how it was used to create a higher level metric to quantify the quality and accuracy of the Hilbert-Huang Transform (HHT) time-frequency representation (TFR).

Finally, as expected we wrap things up with some closing thoughts as well as a variety of future directions and improvements that are now open for pursuit in Chapter 5.

Chapter 2

Background and Prior Work

To lay the groundwork for the need, the ability, and the validity of creating the first of its kind mono-componentness measure are presented in the following sections.

2.1 Terminology

There are a handful of terms whose usage in the time-frequency field are unique. First, is the use of the word “distribution”. When what is now known as the Wigner-Ville distribution (WVD) was first introduced the desire was that it be a proper probability distribution showing the distribution of energy with respect to time and frequency. But it was determined that the WVD actually can, and often does, contain negative values [7]. Obviously, this prohibits it from being a probability distribution. However the term “distribution” continues to be used after decades of prior use in the field.

The second noteworthy terminology “convention” is the interchangeable use of the terms quadratic, bilinear and sequilinear. It is called the quadratic class because the signal is a quadratic term in this class of distributions, rather than a linear term [8, pg. 760]. However, as Janssen explains in [9], they are more precisely bilinear because of the dependence of the energy calculation on two signals. In fact, to border on overly precise, they are sequilinear because the energy is linearly dependent on one signal and conjugate linearly dependent on the second [9].

The final terminological idiosyncrasy is that cross-term and interference term are used interchangeably. But these terms are used to encompass the two unique types of interference, outer interference and inner interference. However, it is the outer interference terms that are so prevalent whether speaking about general interference or specifically outer interference the terms cross-term and interference term are used interchangeably. When speaking specifically about inner interference, the distinction between the two types of interference is made by using the terms “inner” and “outer.” For the remainder of this document cross-term, interference, and interference term are used, as is typical, to indicate outer interference. Only when specifically discussing inner and outer interference will the specific distinction between types be made.

2.2 A Brief Historical Overview

The oldest and most well-known (and studied) member of the quadratic time-frequency class is the Wigner-Ville distribution. In 1932, Eugene Wigner first presented what is known as the Wigner distribution in the field of quantum mechanics in terms of position and momentum [10]. In 1948, Jean-André Ville published his landmark paper on analytic signals that used the analytic signal in the Wigner distribution and in terms of time and frequency [11]. Hence, the Wigner distribution entered the signal processing realm and what is now known as the Wigner-Ville distribution was born.

In time-frequency literature, the names Wigner distribution and Wigner-Ville distribution are used interchangeably. Over the last 15 years or more the name “Wigner-Ville distribution” has become the predominate name used instead of “Wigner distribution.” However, the Wigner and Wigner-Ville distribution are ever so slightly different. The only difference is that given a signal $s(t)$ the Wigner distribution uses it directly, but the Wigner-Ville distribution uses the analytic version of $s(t)$. Overall, the bulk of literature simply uses the name “Wigner-Ville” regardless

of whether the signal being analyzed is analytic or not^{*}. If the analytic signal is used it is specified as so, not assumed simply because the name “Wigner-Ville” was used instead of “Wigner.” Fortunately, in this document and its associated research, which focuses on the multi-component signal interference problem and mono-component signals, the interference geometry is not affected by whether the signal used in the calculation is analytic or not. In actual application one should use the analytic version if the signal is *real-valued*, as this simple step eliminates all cross-terms produced between negative and positive frequencies [14]. This also has the benefit of reducing the required sampling rate of the signal necessary to prevent aliasing in the WVD and WD [2, 14]. Therefore, as it is the norm, in this document we will refer only to the Wigner-Ville distribution regardless and if the analytic signal is used it will be explicitly stated and/or shown as a “pre-cursor” step to the calculation. Now back to the history of the WVD.

The WVD continued to live almost exclusively in the realm of quantum mechanics for several decades after 1948. Along the way, in 1966, Leon Cohen published his article on phase-space quasi-probability distributions, which was focused on quantum mechanics [15]. This particular paper contained a formulation that generalized the Wigner distribution through the addition of a kernel function. This generalization, once adopted by and adapted to signal processing, became known as Cohen’s Class. It is through Cohen’s generalized quadratic joint distribution that all quadratic TF distributions can be represented [16]. In fact, it is through the WVD-plus-kernel formulation that it is seen that all quadratic TF representations (TFRs) are based on the WVD [12, pg. 59][17, 18].

Beginning around 1980 the WVD was re-introduced/re-popularized, along with Cohen’s generalized formulation, through a three article series by Claasen and

^{*}In fact, even one of the world’s leaders in the TF field, Boualem Boashash, sometimes does and does not make the distinction between the Wigner and the Wigner-Ville formulations. For example, in [12, pg. 33], Boashash points out that the WVD refers to the form that uses the analytic signal, but in [13, Chp. 2, pg. 10], he point out that Ville used the analytic signal, but then goes on to indicate in general, analytic signal or not, it is referred to as the WVD. All and all, it is another ambiguous terminology usage of the TF field

Mecklenbräuker published in the Philips Journal of Research [7, 19, 20]. Since that time much of the research in the time-frequency field has focused on the reduction and/or elimination of the cross-terms of the WVD and its quadratic class siblings. These decades of work have not been without breakthroughs. The renewed attention brought by Claasen and Mecklenbräuker works plus Cohen’s Class spurred a burst of research in the early 1980’s. Near the end of the 1980’s attention waned as the new field of wavelets attracted the attention of the community [21, pg. 175]. Then around 1994, there was another significant burst of interest in time-frequency methods and of course the elimination or reduction of cross-terms in the quadratic class. This burst of activity seems to roughly corresponded to publication of several books that are to this day are still highly referenced in the TF field due to their substantial impact on the field such as Cohen’s “Time-Frequency Analysis” [22] and Patrick Flandrin’s “Time-Frequency/Time-Scale Analysis”[†][21]. The research and development in the time-frequency domain continued at an elevated level, compared to the late 1980’s, after the burst of activity around 1994. Roughly, the same level of “output” is still occurring today. Albeit focused on different areas in the time-frequency domain, such as reassignment operations as an attempt to undo the resolution loss caused by the windows and kernel methods used to reduce the impact of cross-terms.

These advancements have been meaningful and provided options beyond just the WVD for generating TFRs with improved resolution (compared to linear methods) and reduced cross-terms, but still suffer from several drawbacks with respect to their usage. These improvements brought the need for more information about the signal under test than is typically available in order to select the “best” distribution for the signal. Still they are all of the quadratic class, with the exception of wavelets, and achieve their improvements through the application of a window/kernel/basis function. The selection of a quadratic class method means a window/kernel/basis must be selected as well as the values for the multiple parameters that typically control said window/kernel/basis function [23]. This again brings the need for

[†]originally published in French in 1993, translated and published in English in 1999

more information about the signal than is typically available. The details regarding selecting which distribution to use and what to set any associated parameters to, plays into a third point which is the tradeoff between how fine a resolution is needed and how much “blurring” the decisions regarding which distribution and how any parameter values will affect the resolution.

It is these challenges that shine a light on why the WVD still holds such appeal. The WVD does not utilize a window, kernel, or basis function [8, pg. 769], [6, pg. 16]. However, to quote the old adage, “You can’t have your cake and eat it too.” The world is full of multi-component signals and the WVD being quadratic results in the generation of numerous interference terms. It is the issue of interference terms in the WVD due to multi-component signals that we wish to examine, but first an introduction to multi-component signals, the Wigner-Ville distribution, and the interference terms that result from the combination of two.

2.3 Multi-component and Mono-component Signals

The definitions of multi-component and mono-component signals are a funny thing. The accepted definition of multi-component signals is a linear summation of mono-component signals, Eq. (2.1), [24, 1], where $x_k(t)$ is the k^{th} mono-component and c_k is its coefficient.

$$x(t) = \sum_{k=1}^N c_k x_k(t) \quad (2.1)$$

There is a definition of a mono-component that was put forward in the early 1980’s, but since then much of the literature that involves multi-component signals makes no attempt to define what a mono-component is or even reference any of the publications that have presented a definition. The handful of papers that do provide a definition of a mono-component signal define it as a signal having an instantaneous frequency (IF) consisting of only one ridge/pattern/IF law at a time or point [1, 25, 26, 27, 28],

[12, pg. 19]. A more elegant wording of what a mono-component is,

“...that is, signals limited to only one time-varying feature in the frequency domain.” — B. Boashash [13]

It is not the most satisfying definition, as in order to determine if the IF forms a single line in a TF distribution one must select a distribution and examine the results. This brings up the perpetual issue of how to choose a TF distribution when little is known about the signal being analyzed and how to tell if it was a good choice given so little information.

There are several points that should be made regarding multi-component signals as a summation.

1. A signal represented as a sum of parts does not mean it is a multi-component signal [29].
2. There are an infinite number of non-unique mono-components which could be summed to create a particular multi-component signal [22, pg. 126].
3. The “original” components composing a multi-component signal are all but unknowable without prior knowledge of how it was created (i.e., the original mono-components used).

Because of the lack of knowledge on how a multi-component signal was created and the specific mono-components used to create it, we need to constrain the possibilities in some other fashion. This has led to the emphasis that mono-components of a decomposed multi-component signal should have some “physical meaning” [29]. While not describable mathematically, it does make sense intuitively.

We live in a physical world. We measure physical events and quantities. Energy, time, and frequency are all physically meaningful units. Time-frequency distributions attempt to localize the energy in a signal to the frequency and time of its occurrence. Consider a musical score. Even a musical piece containing a single instrument

can easily be distinguished as containing different components that are physically meaningful. As an example, the sharp clash of a cymbal and thump of a bass drum all come from the same musical instrument (source) – a drum set – and even if those sounds are produced in isolation or with some duration of overlap, we know intuitively they are components. In fact, even the way in which sheet music is “written” is in “components” localized in time and frequency.

This physical meaning is key with regards to extracting or decomposing a multi-component signal into mono-components. For example, take a pure sine wave with constant amplitude and frequency; this mono-component signal could be represented as a sum of two mono-components with different amplitudes that sum to achieve the original amplitude at any given time, but it is not physically meaningful [30, pg. 66],[24]. There is no advantage or reason to represent a single sine wave as the sum of two signals to represent its amplitude. This same example in a mathematical form is shown below:

$$x(t) = c_1 x_1(t) + c_2 x_2(t)$$

where

$$x_1(t) = x_2(t) = x(t) \quad \text{and} \quad c_1 = 1 - c_2$$

substitution then yields

$$\begin{aligned} x(t) &= (1 - c_2) x_2(t) + c_2 x_2(t) \\ &= x_2(t) - c_2 x_2(t) + c_2 x_2(t) \\ &= x_2(t) \end{aligned}$$


The point to take away from this is that just because a signal can be written as a summation of signals does not make that summation a multi-component signal or the signals summed physically meaningful. Also, even though there are an infinite number of combinations of mono-component signals that could be summed to create a particular multi-component signal and that any of these non-unique sets of mono-components and its likelihood of being the same as the set of mono-components used

in the creation of the multi-component signal is uncertain, it does not create an issue. What matters is that the summation of a set of mono-components equals the original multi-component and that the mono-components are non-trivial with respect to their physical meaning.

2.4 Ridges

The term “ridge” has already been used as how to identify a mono-component signal in a TF plot, but it is important enough concept to merit a brief discussion and to establish the difference between ridges and edges.

Physically, the result of a TF distribution represents energy localized in time and frequency. Ideally, the energy of each mono-component would be concentrated into a knife edge in the TF-plane [12, pg. 30] such as the idealized representations shown in Figure 2.1.



(a) The ideal time-frequency representation of a mono-component signal.

(b) The ideal time-frequency representation of a multi-component signal.

Figure 2.1: In these idealized TF plots each component of a signal is perfectly concentrated to form a single ridge.

In reality, even a mono-component is subject to some flaring at its base similar to how a mountain ridge tapers to its peak. Likewise, in the case of multi-component






Figure 2.2: The “foothills scenario” can be seen in this zoomed in view of the TFR generated via the WVD for a simple mono-component signal.

signals when interference terms are present our ideal mountain ridges are now a mountain range surrounded by foothills. In this analogy, it is the highest mountain ridge that we want to identify while avoiding the “confusion” of the lower ridges of the foothills. An illustration of the “foothills scenario” in the TF plane using the WVD is shown in Figure 2.2.


One point that should be mentioned here is that edge detection and ridge detection are not the same thing. Edge detection in an image seeks the locations where, in the most generic of terms, a discontinuity occurs in a color channel. In the case of the TF representation of a mono-component, edge detection will result in a line being found on both sides of the ridge, Figure 2.3b. Using ridge detection only a line following the peak of the ridge results, Figure 2.3c. At any rate, the concept of what a ridge is in the TF plane is quite intuitive in the context of the topography of a mountain



(a) The WVD results for a linear FM chirp, i.e., a true mono-component.



(b) The Canny edge detection results for the WVD shown in (a).



(c) The ridge detection results for the WVD shown in (a).

Figure 2.3: The TF representation and the edge and ridge detection results for a perfect mono-component signal. Clearly, edged detection and ridge detection produce two very different outcomes.

range.

2.5 The Wigner-Ville Distribution

It has been a bit of a meandering train of thought to get to the point of discussing the Wigner-Ville distribution (WVD), but it is beneficial to frame the discussion of the WVD with the context of multi-component signals, mono-component signals, and ridges.

The WVD is commonly presented in two forms: an auto-WVD and a cross-WVD. On closer examination it can be seen that there is not a significant difference between the two formulations other than notation. It is the signal being analyzed that implies the chosen form/notation. That is, a cross-WVD calculated for a mono-component signal is no different than calculating the auto-WVD of a mono-component. It is simply notational semantics of explicitly labeling the two signal terms as the same signal or different signals. In Eq. (2.2), the continuous form of the auto-WVD is given and from this point on is identified using the notation, $WV_x(t, f)$. Likewise, the cross-WV is shown in Eq. (2.3) and the notation $WV_{x,y}(t, f)$ is used to identify it.

$$WV_x(t, f) = \int_{-\infty}^{+\infty} x\left(t + \frac{\tau}{2}\right) x^*\left(t - \frac{\tau}{2}\right) e^{-j2\pi f\tau} d\tau \quad (2.2)$$

$$WV_{x,y}(t, f) = \int_{-\infty}^{+\infty} x\left(t + \frac{\tau}{2}\right) y^*\left(t - \frac{\tau}{2}\right) e^{-j2\pi f\tau} d\tau \quad (2.3)$$

Equivalently, the auto-WVD and cross-WVD can be calculated in the frequency domain using Eq. (2.4) and Eq. (2.5), where $X(f)$ is the Fourier transform of $x(t)$. It is this formulation that allows for a more efficient computational implementation thanks to the FFT.

$$WV_x(t, f) = \int_{-\infty}^{+\infty} X\left(f + \frac{v}{2}\right) X^*\left(f - \frac{v}{2}\right) e^{-j2\pi vt} dv \quad (2.4)$$

$$WV_{x,y}(t, f) = \int_{-\infty}^{+\infty} X\left(f + \frac{v}{2}\right) Y^*\left(f - \frac{v}{2}\right) e^{-j2\pi vt} dv \quad (2.5)$$

In the time-domain and frequency-domain formulations the term τ and v represent the lag/delay. Often the two inner terms of the WVD, $x\left(t + \frac{\tau}{2}\right) x^*\left(t - \frac{\tau}{2}\right)$ are called the instantaneous auto-correlation function (IAF) or local auto-correlation function (LAF). This terminology also more directly communicates what these two terms and the lag term are doing. They form a symmetrically shifted correlation operation.

The WVD has several nice mathematical properties that make it very appealing. It also has several serious drawbacks to its usefulness. The most important beneficial attribute is its ability to have high resolution in both time and frequency simultaneously. Unlike the short-time Fourier transform (STFT) where there is a direct trade-off between high resolution in one domain at the expense of the resolution in the other domain. Another, key property is that the WVD abides by the quadratic superposition law [30]. This aspect is discussed in Section 2.6. However, this property is also a doubled-edged sword.

It is the fact that the WVD produces a time-frequency representation (TFR) with high resolution in both domains without requiring the selection of a window or other basis type function that make it the ideal TFR. Unfortunately, it is its bilinear use of the signal that makes it nearly unusable due to the prolific presence of cross-terms. Still, it is how the WVD “behaves” when applied to a mono-component signal that is the key to the novel method for determining if a signal is a valid mono-component, to be presented in Chapter 3.

2.5.1 Discrete Wigner-Ville Distribution

There have been several discrete WVD (DWVD) implementations created over the years and much disagreement on which is the best one to use. The problem is multifaceted, but the key main points of disagreement involve aliasing and the preservation of continuous domain WVD properties in the discrete WVD. While the disputes over which is “better” has never been completely resolved, a more or less unspoken decision has emerged that the choice is best left to the requirements of a particular application based on its needs such as, computational speed, sampling rate of the signal, and if the frequency marginal property is needed in the discrete domain.

The frequency marginal property is the continuous WVD property that is especially difficult to retain upon transition into the discrete domain. In short, the frequency marginal property of the continuous WVD amounts to the fact that in the continuous domain the spectral energy density (a.k.a. energy spectrum) can be calculated by integrating the WVD across time, as shown in Eq. (2.6) where $S(f)$ is the Fourier transform of $s(t)$.

$$\int WVD_s(t, f) dt = |S(f)|^2 \quad (2.6)$$

As a discussion of the many additional mathematical properties possessed by the WVD and the effects going from a continuous domain to a discrete domain is well beyond the scope of this work, the author recommends that those interested in this information see [7, 31, 32]. The pertinent detail for this work is that the frequency marginal property has traditionally been lost when going from the WVD to the DWVD. That is until the work of O’Toole that resulted in an improved discrete analytic signal transform and a new DWVD implementation [31].

Ultimately, for our purposes, as presented in Chapter 3, it is the computational speed and aliasing attributes that are important. The aliasing as expected is dependent on the sampling rate. The computational speed is tied to the sampling grid used to go from the continuous domain WVD to the DWVD.

By using the analytic version of real-valued signals not only are all cross-terms between negative and positive frequencies eliminated but it also reduces the sampling rate required to prevent aliasing from two times the Nyquist sampling rate (four times the highest frequency we wish to retain) back to the typical Nyquist rate [14].

Using the analytic signal (of a real-valued signal) also has the advantage of reducing the aliasing by ensuring that the energy contained in the negative frequencies is zero and thus does not overlap with positive frequencies [14]. Except, it turns out that using the traditional methods of calculating the analytic version of a real signal does not actually result in all of the negative frequencies being nulled. There is a new method for generating the analytic signal that comes very close to truly zeroing all of the energy in the negative frequencies. This method was also developed by O’Toole along with his DWVD implementation which retains the most WVD properties of any discrete implementation created to date [18, 31, 32, 33]. Here again the creation and proofs of both the improved analytic signal generation algorithm and the O’Toole DWVD implementation are beyond the scope of this work. We simply make use of these advances in order to produce the pioneering work in Chapter 3. A fast and memory efficient versions of these algorithms are provided by O’Toole and are utilized to ensure the correctness of the analytic signal generation and DWVD calculation [34, 35, 36, 37, 38].

2.6 Interference Terms in the Wigner-Ville Distribution

With the basics of the WVD, mono-component, and multi-component signals presented we can now address the issue of interest — the interference terms produced in the WVD when it is applied to a multi-component signal. This issue is extensive and detrimental to the use and interpretation, by both humans and machines, of TF representations produced of multi-component signals using the WVD (or any

quadratic TF distribution for that matter) [4, 39].

During the time-frequency rebirth of the early 1980's, extensive research was done focusing on the source of the cross-terms, where they come from, where they will occur, etc. It was found that there are two types of cross-terms, which became known as outer interference and inner interference. Outer interference only occurs between the mono-components composing a multi-component signal. This is the most prevalent contributor to the cross-term problem. Inner interference occurs in the concaves of the instantaneous frequency (IF) of a single component [24]. For example, inner interference would be seen in the WVD results for a mono-component signal having a sinusoidal IF. Inner interference can be present in multi-component signals as well. In these cases, the inner interference would be associated with the IFs of any mono-components that are part of the multi-component signal and contain concaves. To state it more generally, inner interference occurs in the WVD for mono-components with non-linear frequency modulation [40, pg. 231][41].

It is the outer interference terms which are the focus of this work. They are the most prolific and occur in *all* multi-component signals. Whereas, not all multi-component signals contain one or more mono-components with IFs that have concaves in them, but *all* multi-component signals will contain outer interference [30].

The effect a multi-component signal has on the WVD can easily be seen in the following example, which uses a simple two component signal and Eq. (2.2).

Let $s(t)$ be a multi-component signal resulting from the sum of two mono-component signals, $c_1 s_1(t)$ and $c_2 s_2(t)$, such that $s_1(t) \neq s_2(t)$ and where c_1 and

c_2 are some arbitrary coefficients [30].

$$s(t) = c_1 s_1(t) + c_2 s_2(t) \quad (2.7)$$

$$\begin{aligned} WV_s(t, f) &= \sum_{k=1}^2 \sum_{l=1}^2 c_k c_l^* WV_{s_k, s_l}(t, f) \\ &= c_1 c_1^* WV_{s_1, s_1}(t, f) + c_2 c_1^* WV_{s_2, s_1}(t, f) \\ &\quad + c_1 c_2^* WV_{s_1, s_2}(t, f) + c_2 c_2^* WV_{s_2, s_2}(t, f) \end{aligned} \quad (2.8)$$

Expanding the two cross-WVD terms in Eq. (2.8) we have

$$WV_{s_1, s_2} = c_1 c_2^* \int_{-\infty}^{+\infty} s_1\left(t + \frac{\tau}{2}\right) s_2^*\left(t - \frac{\tau}{2}\right) e^{-j2\pi f\tau} d\tau \quad (2.9)$$

and

$$WV_{s_2, s_1} = c_2 c_1^* \int_{-\infty}^{+\infty} s_2\left(t + \frac{\tau}{2}\right) s_1^*\left(t - \frac{\tau}{2}\right) e^{-j2\pi f\tau} d\tau \quad (2.10)$$

Recognizing that the s_1 and s_2 terms in Eq. (2.9) and Eq. (2.10) each forms a cross-correlation but with the shift being split between the two terms. We know that the definition of cross-correlation between two sequences is:

$$R(\tau) = \int_{-\infty}^{+\infty} x(\tau) y(t + \tau) d\tau$$

Using a dummy variable, β , and letting $\beta = \tau + t$, and $d\beta = d\tau$ we see

$$R(\tau) = \int_{-\infty}^{+\infty} x(\beta - t) y(\beta) d\beta$$

Since β is a dummy variable, a lag/delay variable, just like τ we simply revert to the original dummy variable notation leaving us with

$$R(\tau) = \int_{-\infty}^{+\infty} x(\tau) y(t + \tau) d\tau = \int_{-\infty}^{+\infty} x(\tau - t) y(\tau) d\tau$$

It is through this and the recognition that the two cross-WVD terms are each a cross-correlation function and the Fourier transform of the cross-correlation function at that, which then allows us to say that:

$$WV_{s_1, s_2}^* = WV_{s_2, s_1}$$

yielding

$$\begin{aligned} WV_s(t, f) &= c_1 c_1^* WV_{s_1, s_1}(t, f) + c_2 c_2^* WV_{s_2, s_2}(t, f) \\ &\quad + c_1 c_2^* WV_{s_1, s_2}(t, f) + c_1^* c_2 WV_{s_1, s_2}^*(t, f) \\ &= |c_1|^2 WV_{s_1} + |c_2|^2 WV_{s_2} + 2 \operatorname{Re} \left\{ c_1 c_2^* WV_{s_1, s_2} \right\} \end{aligned} \quad (2.11)$$

where WV_{s_1} and WV_{s_2} are the auto-WVDs of the s_1 and s_2 components, respectively and WV_{s_1, s_2} is the cross-WVD of the two components. From this it is easy to see that for the signal $s(t) = s_1(t) + s_2(t)$ the cross-term of $WV_s(t, f)$ is contributed by $2 \operatorname{Re} \{ c_1 c_2^* WV_{s_1, s_2} \}$. The cross-term is two times the real-part of the cross-WVD of the two generating components. This means that it is entirely possible for the cross-term to have an amplitude that could be as much as twice the product of the amplitudes of the two components that produced the cross-term [42, 43]! How close to the maximum possible amplitude a cross-term possesses depends of course on the signs of the two components. The closer to parallel the two generating components are the closer to the worst case cross-term amplitude it will be as well [1]. This example also serves to illustrate that the WVD holds to quadratic superposition.




Figure 2.4: The time, frequency spectrum, and TF representation (Wigner-Ville) are shown of two FM burst with Gaussian amplitude envelopes. The outer interference is clearly shown lying at the mid-points between the two atoms. The lower left Gaussian atom is centered at $(t, f) = (38, 0.1)$, the upper right atom at $(96, 0.35)$.

We can see the effects of Eq. (2.11) in Figure 2.4 which shows the WVD results for a two-component signal composed of two atoms each with a constant frequency and a Gaussian amplitude. The analytic version of the signal is used in this example. Most importantly, notice the cross-term located at the mid-point between the centers of the two atoms. This cross-term is specifically outer interference.

In fact, the reason it could be mathematically shown in Eq. 2.11 that the WVD of a multi-component signal is not equal to the sum of the WVD of each mono-component is because the mono-components, $s_1(t)$ and $s_2(t)$, were known *a priori*.

While the less pressing issue of inner interference will not be dealt with explicitly

in this work, due to its lower impact in comparison to the guaranteed presence of outer interference terms in *every* multi-component signal, it does merit at least an illustration as a point of reference. Figure 2.5 shows a mono-component signal with the sinusoidal frequency modulation represented by Eq. (2.12).

$$x(t) = A_c \cos \left(\omega_c t + \frac{\Delta A_m}{\omega_m} \left(\sin \left(\omega_m t - \frac{\pi}{2} \right) + 1 \right) \right) \quad (2.12)$$

where

- A_c is the amplitude of the carrier.
- ω_c is the angular frequency of the carrier
- Δ is the maximum frequency deviation
- A_m is the amplitude of the modulating sinusoid
- ω_m is the angular frequency of the modulation

The presence of inner interference is clearly visible in the concaves. In fact, in the case of the WVD it has been determined that inner interference is only present in mono-components with non-linear frequency modulation (FM), i.e., signals whose IF is the result of a phase function of order two or higher [12, Chp. 3]. Fortunately, the developments presented in Chapter 3 provide a basis that is extensible to also allow a methodology for detecting and addressing inner interference. See Section 5.2.

To elaborate a bit more on the geometric structure of interference terms it has been shown that the interference will occur at the midpoint of any linear cord that can be drawn between two non-zero (t, f) points on the instantaneous frequency law and will oscillate. The oscillating frequency increases as the distance between the two generating points increases. This behavior has been exhaustively derived and proved in the definitive sources on this behavior: [44, 24, 45, 24, 30].

Returning to the WVD equation for the two component example, Eq. (2.11), and expanding upon this using the definition of an N -component signal defined earlier, Eq. (2.1), the WVD, separated into auto-terms and cross-terms, for an N -component signal can be derived [30]. This is all well and good until one recalls that




Figure 2.5: The WVD of a sinusoidal frequency modulated signal and serves to illustrate inner interference terms. The exact IF law is the bold red line plotted on top of the WVD.

the components that make up a multi-component signal are almost never known.

$$WV_s(t, f) = \sum_{k=1}^N |c_k|^2 WV_{s_k} + 2 \operatorname{Re} \left\{ \sum_{\substack{k=1 \\ k \neq l}}^N \sum_{l=1}^N c_k c_l^* WV_{s_k, s_l} \right\} \quad (2.13)$$

The fact that the WVD possesses the key property of conforming to quadratic superposition is visible in Eq. (2.13). This means that for every unique pair of components there will be two auto-terms and a cross-term. The more components in a multi-component signal the more cross-terms you have. The number of cross-terms grows at a quadratic rate with respect to the number of components. If the number of components, N , is known beforehand it has been shown that there would be N auto-terms and the number of cross-terms could be calculated using Eq. (2.14) [30].

$$\left(\frac{N}{2} \right) = \frac{N(N-1)}{2} \quad (2.14)$$

The magnitude of the cross-term problem is revealed now that we have seen that,

- cross-terms grow quadratically with linear component growth,
- cross-terms can have amplitudes twice the product of the amplitudes of its generating components,
- and consider the fact that the cross-terms can just as easily occur on top of an auto-term.

It is because of the WVD's interference terms that the usability, readability, and overall usefulness for signal analysis is limited [46][12, pg. 32]. The interference terms can be removed through the use of an appropriate kernel, but the resolution will be reduced [47] and to keep this reduced resolution as high as possible requires that the most appropriate kernel be used. The selection of this kernel requiring *a priori* knowledge of the multi-component signal which is seldom if ever available until after

using TF analysis.

2.7 Summary

It is these key foundation stones, multi-component and mono-component signals, the WVD, the resolution of the WVD versus the issue of interference terms, and how the interference terms occur, that the following work builds upon. They also provide the context for such decisions as why the only TF distribution used is the WVD and why the result when the WVD is applied to a mono-component is important. It is through the knowledge and context of these fundamental elements of time-frequency analysis, multi-component signals, and interference structure that opened the door to a first of its kind development.

By focusing on the mono-components of a multi-component, instead of cross-terms of a multi-component, the auto-terms (and thus mono-components) are now of first importance. The key to all of this is to be able to determine if a signal is a true mono-component or not. It is through this change of perspective and that the realization that an algorithmic, objective determination of a signal's mono-componentness is the keystone to allowing a broader, more practical use of the WVD.

Chapter 3

The Measurement of Mono-Componentness

For decades, the approach for dealing with the cross-terms of a multi-component signal has been to directly attack the problem mathematically by “adapting” the WVD through the addition of different kernels [22, Chp. 9]. If we too focus strictly on trying to eliminate cross-terms, we will be stuck in this same situation that has existed since the interference geometry of the WVD was determined [24, 1]. That path has been investigated for decades and has yielded valuable knowledge about designing time-frequency kernels with improved cross-term suppression that reduces or eliminates the interference terms of a specific signal, but always requiring *a priori* information about the signal and sacrificing resolution.

Significant work has been done by Hlawatsch and Flandrin, [24, 30], in which the geometric attributes of the interference terms, along with closed form equations for the independent calculation of the auto-terms and cross-terms, of the WVD were derived with respect to multi-component signals. This landmark research came with one critical assumption — that the mono-components forming the multi-component signal are known beforehand. Still, the relationship between each unique pair of components and the effects the locations and distance between the two components

in the TF plane have on the interference term produced by the pairing are known. However, Hlawatsch and Flandrin’s work has been treated largely as informational — highly cited, often reproduced in reference books, useful information, but not necessarily actionable information given the monumental requirement of having access to **all** the mono-component signals that when combined form the multi-component signal being analyzed. However, this information is the foundation of my method for reliably and objectively determining if a signal is mono-component or not.

Instead of focusing on the interference terms, as has been the case until now, when the unique approach of focusing on mono-componentness was hatched. Specifically, a means for validating the mono-componentness of a signal. This is because if a set of *legitimate* mono-components which sum to equal the original multi-component signal can be determined then the interference issue can be approached from an entirely different perspective.

The approach presented in the following sections does not seek a “magic” generalized kernel. Nonetheless, it eliminates interference terms in the WVD results for any multi-component signal without requiring any significant prior knowledge or years of expertise. Realistically and practically speaking, there is not much hope for the existence of a perfect generalized form that works for everything. Consider that the recording of a piano solo, an opera, and an orthogonal frequency-division multiplex (OFDM) cellular transmission are all one-dimensional multi-component signals. The diversity of content and composition that is encapsulated in the humble phrase, multi-component signal, makes the successful discovery of a perfectly generalized kernel likely impossible.

Therefore, I present a unique, different approach that leverages the experience and wisdom of over a half century of time-frequency research.* Instead of attacking the problem through modifications and adaptation of the WVD, we should focus on the things we know and have with certainty.

*Wisdom is knowledge and experience is the application of wisdom. Likewise as the colloquialism goes, “Insanity is doing something over and over again and expecting a different result.” You may not be insane, but you certainly are going to be frustrated beyond measure.

The things we have and know with confidence are:

1. The signal under test.
2. That a multi-component signal can be represented by a linear summation of mono-components.
3. The WVD obeys the quadratic superposition principle.
4. That *if* all the mono-components are known then the auto-term of each mono-component and the cross-term from each unique pairing of mono-components can be calculated independently.

It is not a very long list and its contents do not invoke a lot of awe. However, if we return to Eq. (2.11) of the two component example in Section 2.6, but approach it from the opposite point of view by examining the behavior of the WVD when a single component signal is used, an additional item can be added to this list. One of great significance and is the keystone of the first-of-its-kind quantitative measure that will now be presented.

3.1 Wigner-Ville distribution of a Mono-component

Earlier, in Section 2.6, the WVD of a two component signal was calculated to show mathematically where the cross-terms originate. However, what if we were to consider the problem from a different perspective? What happens if the WVD is calculated of a signal that turns out to be a mono-component?

Defining a single mono-component signal, $h(t) = \frac{1}{2}h_1(t) + \frac{1}{2}h_1(t)$, where the single component has been split into two identical pieces each possessing half of the original amplitude. This decomposition is physically meaningless as in reality adding two completely identical signals results in a single signal whose amplitude is the sum of the parts. Next, by performing the same expansion, as used in Eq. (2.11), upon $h(t)$ yields Eq. (3.1).

$$\begin{aligned}
WV_h(t, f) &= \left(\frac{1}{2}\right)^2 WV_{h_1, h_1}(t, f) + \left(\frac{1}{2}\right)^2 WV_{h_1, h_1}(t, f) \\
&\quad + \left(\frac{1}{2}\right)^2 WV_{h_1, h_1}(t, f) + \left(\frac{1}{2}\right)^2 WV_{h_1, h_1}^*(t, f) \\
&= \frac{1}{4}WV_{h_1} + \frac{1}{4}WV_{h_1} + \frac{1}{4}WV_{h_1} + \frac{1}{4}WV_{h_1} \\
&= WV_{h_1}(t, f) = WV_h(t, f), \text{ since } h = h_1
\end{aligned} \tag{3.1}$$

This clearly shows that the result of the WVD for a mono-component is a single term, which affirms the definition of a mono-component as a single ridge in the time-frequency domain. There is the special case of inner interference which results in the WVD of any valid mono-component with a non-linear phase function/instantaneous frequency [41] producing more than one ridge, but this is addressable and will be discussed in Section 5.2.1. Regardless, herein lies the key to eliminating outer interference terms from the WVD's TFR of a multi-component signals or any QTFD that has the quadratic superposition property for that matter. That key is the ability to determine if a signal is a multi-component or mono-component signal and the basis to do so with confidence. For if the TFR produced via the WVD has only a single ridge then it must be a mono-component!

Consider this, if the WVD for a single (t, f) point of an unknown multi-component signal is calculated the WVD results will contain $N + \frac{N(N-1)}{2}$ corresponding locations, where N is the number of components in the multi-component signal (a detail not available to us). From Eq. (2.11) we see that for any single time-frequency pair (t, f) , in the two-component signal ($N = 2$), Eq. (2.7), results in three corresponding (t, f) locations in the WVD result – two for the auto-terms and one cross-term. However, despite the trivialness of the example shown in Eq. (3.1), the calculation of the WVD for a single (t, f) point of a mono-component signal results in a single (t, f) location in the resulting TFR. This is a fact we can rely on. A property that results in a trait we can detect that only occurs in one specific situation — the WVD was calculated

for a true mono-component!

This is the critical, deterministic attribute that allows mono-componentness to be detected, despite the fact that the only definition of a mono-component signal is more conceptual than mathematical. This is a unique attribute of the WVD that only occurs for one type of signal. It is not signal dependent, kernel dependent, and does not even involve a single parameter that must be configured. It also requires **no** prior knowledge of the signal.

This beautiful, reliable, and intrinsic behavior of the WVD adds an additional, extremely beneficial, item to our list of things we have and know with certainty.

1. The signal under test.
2. That a multi-component signal can be represented by a linear summation of mono-components.
3. The WVD obeys the quadratic superposition principle.
4. That *if* all the mono-components are known then the auto-term of each mono-component and the cross-term from each unique pairing of mono-components can be calculated independently.

5. The WVD of a mono-component forms a single ridge in the resulting TFR[†]

The challenge of constructing a measure of mono-componentness has been to ensure the validity of such a measure and the measuring of something with such an abstract definition; albeit an accepted definition/description [22, 48]. Having presented the foundation that provides both legitimacy and the hints of how a mono-componentness can be “measured” (i.e., a true mono-component’s WVD has a single ridge), we can now delve into how to actually “measure” mono-componentness algorithmically.

[†]The exception being the special case of mono-components with polynomial phase functions.

3.2 The Algorithm

Now, that a means for uniquely identifying a mono-component versus a multi-component has been identified we can delve into how mono-componentness is measured algorithmically based on what we now know. The actual means for creating an objective method for determining the mono-componentness of a signal is elegant in its simplicity. The key steps are:

1. Calculate the Wigner-Ville distribution.
2. Detect the ridges.
3. Count the ridges.

Figure 3.1 shows a flowchart of the algorithmic implementation developed to calculate the mono-componentness of a signal.

There are several practical issues facing us in transitioning from knowing that the WVD of a mono-component is a single ridge to actually determining this in the blind and with minimal user interaction, i.e., parameter tweaking. Issues such as, the WVD results of a multi-component signal can easily, and typically does, contain a multitude of ridge fragments due to the oscillating amplitude of cross-terms that the ridge detector will attempt to link into ridges even though they are superfluous. Then there is the problematic situation where a cross-term and auto-term directly overlap one another, which causes the negative values of the cross-terms to fragment what should be the continuous ridge of the auto-term into an unknown number of small segments obscuring the relationship between those segments. On top of that many ridge detection algorithms require prior knowledge of the number of ridges to be extracted and will extract the requested number of ridges regardless of whether they exist or not [49].

But before we focus on these implementation issues let us first examine the key steps which make up the actual implementation. Having previously discussed




Figure 3.1: Flowchart of mono-componentness measure.

the WVD (Section 2.5), the interference terms produced when analyzing a multi-component signal using the WVD (Section 2.6), and the effects of applying the WVD to a true mono-component (Section 3.1), we will proceed directly to the next steps in the algorithm – ridge detection and ridge counting.

3.3 Ridge Detection


Ridge detection is a key element of reliably determining if a TFR is showing a mono-component or a multi-component. In fact, even though there are very few attempts at defining what a mono-component is they all eventually condense to one general point. That a mono-component is a signal whose TFR exhibits a single ridge — the definition of a mono-component [25, 27]. Clearly, a stable and reliable ridge detection is critical to objectively quantifying the mono-componentness of a signal.

Often ridge detection and edge detection are used as if they are one and the same. This is not true. Especially for this application where many edges exist naturally due to the discrete implementation of the WVD.


The detection of ridges has often been used as a means for determining the instantaneous frequency (IF) of a signal from its TFR [50, 51]. Ideally a mono-component would be perfectly concentrated along its IF to form a ridge. In reality, signals acquired from the real-world are noisy, contain quantization effects, and limitations relating to the fact that a discretized TFR algorithm is used. All of these items ultimately result in the dispersion of a ridge. Using an edge detection algorithm, such as the Canny edge detector, results in identifying the boundaries of the ridge structure at best. Because of the discretized frequency and time “bins” along with the oscillatory tendencies of the WVD, edge detection, even for a known synthetic mono-component test signal, results in numerous edges being detected as shown in Figure 3.2b. By using ridge detection the desired information is targeted — the center of mass, the skeleton of the ridge structure instead of all the edges as can be seen in Figure 3.2c.

The topography of a mountain range is an apt analogy for edges versus ridges. Detecting all the edges is equivalent to identifying all of the crests of the foothills and lower peaks that abut the tallest mountain range. Ridge detection is kin to identifying just the tallest stretch of mountains while ignoring all the foothills and lower ranges.


In order to be consistent with our goal of requiring no prior knowledge of the



(a) The WVD results of a simple, 2-component multi-component signal.



(b) The Canny edge detection results for the WVD shown in (a).



(c) The ridge detection results for the WVD shown in (a).

Figure 3.2: The WVD TF representation, Canny edge detections using the edge sensitivity thresholds $[0.3, 0.6]$, and Crazy Climbers ridge detection results in a simple, 2-component multi-component signal consisting of two parallel linear FM chirps.

multi-component being analyzed requires that any ridge detection algorithm used be capable of detecting multiple ridges in a TFR without being told how many ridges exist before hand. While this might seem an intuitively simple requirement it proved to be the deciding factor. For as it turns out many existing ridge detection algorithms require the number of ridges sought *a priori* [49]. However, there does exist a well suited ridge detection algorithm for the detection of multiple ridges in a single TFR without any prior information – the Crazy Climbers algorithm.

3.3.1 Crazy Climbers Algorithm

The Crazy Climbers algorithm was introduced in [52]. At its core it uses simulated annealing but for a fundamentally different purpose. Instead of seeking to minimize an objective/penalty function it uses a simulated annealing styled procedure of temperature cooling but with no constraints on the path. That is to say, finding the global or near global maximum value in the TF plane does not terminate the procedure. Instead, the “climbers” seek to find higher ground, but once a high point (local maxima) is reached instead of staying there they will seek to move to a neighboring high point. Thus the climbers leverage the maximum seeking behavior of simulated annealing but without the objective of finding the global maximum. The climber is free to move in any direction, but prefers (thanks to the simulated annealing basis) to move to a neighboring point that is part of the energy concentration forming the ridge.

Another key aspect of the Crazy Climbers (CC) algorithm is that “visits” by the climbers are recorded forming an occupancy grid. CC does not return the value of the local maxima. It is actually compiling an occupancy measure. Many climbers, all starting from random initial locations on the TF grid, seek higher ground and want to follow high points (high with respect to each points neighboring points) which connect (abut) to other “high” points. As these climbers travel around the TF grid seeking ridges the number of “visits” each location on the grid receives is accumulated


and produces an occupancy measure.

One last difference worth mentioning between simulated annealing and the Crazy Climbers algorithm is that the domain in which Crazy Climbers is applied is quite small, just two discrete degrees of freedom.

Since for a multi-component signal at any given (discretized) instant in time their exist more than one frequency on the TF plane (otherwise it would be a mono-component) each climber will first look to move along the time axis. The climber must take a step along the time axis. The rationale for this small but important detail is elegantly clear when viewed from a practical stand point. Time stops for no one. Even if the frequency is constant, time is still ticking along. In CC, once a climber has taken a step along the time axis it will then evaluate the two possible steps along the frequency axis (an increase in frequency or decrease). However, unlike the time axis move the climber is not required to make a move along the frequency axis. Instead, the frequency axis move is evaluated in the normal simulated annealing style by evaluating the probability of the move which is dependent on the current temperature [52, Chp. 7]. In the end, as climbers find peaks and attempt to move along a path that consists of steps to the largest value it can without going backward in time, even if the step is to a value lower than its current position each position accumulates visits by climbers. The number of visits for each position in the TF plane is tallied ending in an occupancy measure. Figure 3.3a shows the two-dimensional input, a TFR, and with the occupancy results, Figure 3.3b, produced using CC. The next step is to construct ridges from the occupancy measures through a process called chaining.

3.3.2 Chaining

The next step of the Crazy Climbers algorithm is to take the occupancy measure and “chain” these values together to extract the ridges. While the climber’s paths seek the high, abutting values over lower ones they are just time-frequency grid locations



(a) The WVD results as supplied as the input to the Crazy Climbers ridge detector. (b) The occupancy results for the WVD shown in (a).

Figure 3.3: The TFR and the occupancy grid produced by the first stage of the Crazy Climbers algorithm i.e., prior to the chaining stage.

until chained together to form curves/ridges. This process is the key to extracting ridges. The occupancy information locates local maxima in the 2-D time-frequency plane. The chaining process of the Crazy Climbers algorithm has two steps: 1) Threshold the occupancy scores. 2) Propagate (“chain”) a path, if possible, from the remaining occupancy values by seeking the highest occupancy value within the step-size[‡] neighborhood of the current location [52, 51].

As the highest occupancy points are connected a ridge is created. At the point when there are no longer any locations (which remained after thresholding) that can be reached by taking a step the ridge is considered to have ended.

However, between the step-size, the fact that climbers must move in time but not necessarily in frequency, and that the occupancy measurement of Crazy Climbers is calculated upon a discretized representation of the TF plane, can lead to what should be a single continuous ridge being erroneously segmented by single cell (pixel) gaps.

[‡]The step-size taken is a user supplied parameter.

3.3.3 Gap Filling

Small gaps commonly occur in what should be a continuous ridge. These gaps occur when small (typically one frequency “bin”) gaps occur in a ridge’s energy. These gaps are due to the discretization of the domain. The gaps propagate into the Crazy Climbers occupancy grid because a climber is never forced to make a move in the frequency direction during each of its steps. Since the climbers are seeking neighboring, connected “high” values, the shift due to discretization is not overcome because the climber takes a step in time, but no step in frequency, resulting in its isolation in the next iteration because there are no other “high” abutting point to its current location and it cannot travel backwards in time, thus ending the chain and ridge. In order to escape the discretization gap it would require a two “bin”/“cell” step in one domain (frequency) but only a single bin step in the other domain (time), which is not allowed by the chaining algorithm as gaps in frequency are the key attribute used to distinguish different ridges.

While this situation sounds like a significant issue it is not, the author has developed a solution that leverages the fact that the chained occupancy measure results can be treated as an image of labeled blobs. Any single pixel gaps can then be filled using morphological bridging from the image processing domain with a structured element known to fill one pixel gaps in lines. Morphological bridging is simply the process of performing a morphological dilation using a structural element consisting of a three-by-three block of all ones. Next, the dilated results undergo a morphological erosion operation using the same structural element. This results in the bridging of all single pixel gaps between the labeled blobs. By leveraging the Crazy Climbers algorithm’s chaining process, along with the morphological operations from the image processing domain, a simple, effective gap bridging and thus ridge merging procedure was developed.

3.4 Implementation Considerations

The selection of the Crazy Climbers ridge detection algorithm freed us from the issue of needing to know how many ridges are present in the WVD’s TFR *a priori*. It also provides the number of ridges detected as a naturally occurring output from its chaining stage. This is just one example of how by breaking up the issues that arise in actually implementing a mono-componentness measure, we can structure the algorithm to make use of every available advantage, even using characteristics traditionally considered as detrimental to our benefit.

As just mentioned, the CC ridge detector addresses the challenge of blindly determining how many ridges are present, but there is still the point of concern involving the overwhelming number of ridges, especially small fragments and the associating of fragmented ridge pieces into a single ridge.

3.4.1 Auto-WVD or Cross-WVD?

The first question that must be answered is whether the cross-WVD or the auto-WVD equation should be used. In this case, we take advantage of the quirk that the difference between the auto-WVD and the cross-WVD is purely notational unless two different signals are used, see Section 2.5. Since we are measuring of the mono-componentness of a signal there is only ever a single signal presented to the measure. As such there is not actually a choice to be made; the auto-WVD and cross-WVD of a single signal are identical.

3.4.2 Sensitivity and False Positives

Another point of concern, that is anything but trivial, is the problem of false positives. In this work, a false positive occurs when a signal is declared to be a mono-component, but is actually a multi-component signal. The potential for false positives to occur happens when there is a dominant mono-component and one or more significantly

weaker mono-components present.

The issue is two-fold. One being the fact that TFRs are traditionally analyzed visually as if they were an image. The second being the potential for a wide dynamic range of amplitudes compounded by representing this range using a color map and that a numerically different color does not mean it is perceptually different. The end effect is that low amplitudes are quickly lost due to the “brightness” of the dominant amplitude given the dynamic range that must be represented (and distinguished) by color alone. Thankfully, by using a ridge detection algorithm we are able to detect ridges that are not perceptible to the human eye. Still, how does this help mitigate false positives? Since, we are already in the vein of examining things from a different perspective we will not stop now.

One of the major contributors to the problem of interpreting and using TFRs generated using the WVD is that the cross-terms are not only prolific, but have the potential to be twice the amplitude of the contributing pair of mono-component, Section 2.6. This doubled amplitude situation is the worst case scenario, but still in all situations the amplitude of both components contributes to the amplitude of the cross-term. Needless to say, this behavior is a serious problem when visualizing the TFR results of the WVD, but for determining mono-componentness it is actually an advantage. It enables the presence of multiple components to be detected even if one component is dominating the other(s), through the presence of the cross-terms resulting from $2 \operatorname{Re} \{c_1 c_2^* W V_{s_1, s_2}\}$. This means that despite the weakness of one component compared to the other both will contribute to the cross-term’s amplitude allowing the cross-term to continue to be detectable by the ridge detector.

This cross-term amplitude behavior is most easily understood visually. In Figure 3.4(a) and (b), a 2-component signal composed of two frequency modulated tones with identical amplitudes are shown in the joint TF domain, Figure 3.4a, as well as in a tri-plot showing the individual time and frequency domain plots perfectly aligned with the joint TF domain plot, Figure 3.4b. As a brief aside, notice in the frequency domain subplot in Figure 3.4b that there is no corresponding

frequency content for the cross-term. This drives home the fact that cross-terms are a mathematical artifact of the bilinear formulation of the WVD and have no physical aspect.

In Figure 3.4(c) and (d), the same 2-component signal is used, but in this case one component has an amplitude that is a tenth of the other component. Even by eye, you can see that while only one of the two auto-terms is visible the cross-term is also still visible.


A powerful example of this can be seen by examining the frequently used TF test signal consisting of a single echolocation pulse produced by a Large Brown Bat, *Eptesicus Fuscus*, that has been used extensively[§]. Just a few of the most well known references that use this signal are: [23, 53, 54]. Often this signal is noted as containing three components [23, 53] [54, Chp. I, Chp. 5]. However, it can be clearly shown that it must be a 4-component signal based solely on the cross-terms.

By carefully examining the TFR, an oscillating cross-term can be seen above the “third” auto-term. The WVD’s interference geometry [30] shows that in order for this cross-term to exist it must occur between two auto-terms.

The key is noticing that there is a cross-term (identifiable as a cross-term by its oscillation) that is not between any of obvious components. By stretching the color map to improve the discernibility between small amplitudes, at the expense of the discernibility between the stronger amplitudes, we can see the very weak fourth auto-term that had to exist in order for the “extra” cross-term to occur between two components.


It is only through what has long been considered a serious disadvantage to using the WVD that we are able to have confidence that the mono-component measure will almost never produce a false positive. While a means for proving the probability of false positives is problematic due to the vast diversity of multi-component signal, the cross-term amplitude is signal independent providing the confidence that there will

[§]This signal and the permission to use it was graciously provided by Curtis Condon, Ken White, and Al Feng of the Beckman Institute of the University of Illinois and is publicly available at <http://dsp.rice.edu/software/bat-echolocation-chirp>



(a) A 3-D plot illustrating amplitude difference between cross-terms and auto-terms of the WVD results of a simple, 2-component multi-component signal where the components are equal in amplitude.


(b) The time, frequency, and 2-D TF domain plots of the same signal shown in 3.4a. The spectral content shown in the frequency domain plot clearly shows that the oscillating cross-term is not physically present in the signal, but an artifact of the WVD calculation.



(c) The 3-D plot of the WVD results of the same 2-component signal as in 3.4a, but this time one of the components has an amplitude one tenth the amplitude of the other.

(d) The time, frequency, and 2-D TF domain plots of the same signal shown in 3.4c. The spectral content plot does not even show a second component because of the dominant tone, but the presence of the cross-term makes it clear there are two components present.

Figure 3.4: 2-D and 3-D plots showing how the cross-terms are detectable thanks to their elevated amplitude even if one of the contributing mono-components has a low amplitude.



(a) The WVD of the classic bat echolocation chirp TF test signal. Notice the “extra” cross-term (circled) which requires a fourth auto-term to exist above it.

(b) The same WVD of the bat echolocation chirp, but with the color map augmented to accentuate low amplitudes the very weak fourth auto-term is visible (circled) above the “extra” cross-term.

Figure 3.5: A real-world illustration of how the detectability of the cross-terms, due to their higher amplitude, provides information on the existence of an all but imperceptible auto-term.

likely never be a false declaration of mono-componentness.

In fact, in testing it has been repeatedly determined that without the use of a minimum required amplitude (elevation) threshold in the ridge detector, an overwhelming number of superfluous ridges would result from chaining together near zero occupancy values. This is true even when a perfect synthetic mono-component signal is being tested. In the end, the elevated amplitude of the cross-terms allows a minimum occupancy threshold to be used to reduce the near-zero “clutter” with little risk of causing a false positive. Ultimately, the elevated amplitudes of cross-terms all but guarantee there will be no false positives. In fact, we go from worrying about a lack of sensitivity to actually having a propensity to being over-sensitive.

3.4.3 Handling Large Numbers of Ridge Detections

Finally, there is the issue of detecting, connecting, and the handling of the WVD results containing a multitude of ridges regardless of their sources. The fact that falsely declaring a signal to be a mono-component is not a concern thanks to the intrinsic trait of the WVD is an advantage, as discussed in Section 3.4.2. On the other hand, it also means that there is likely to be many ridge detections. Given the information in Section 3.4.2 it would appear that the challenge of this issue has only increased in difficulty.

This did appear to be the case at first. Remember, that golden gem of knowledge regarding the one reliable thing that allows the objective determination of mono-componentness? That the WVD results of a true mono-component will contain only one ridge. Well, that is again our saving grace for dealing with prolific numbers of ridges. The key is focusing on what we want to know, “Is the signal being tested a mono-component?” How do we know? It has a single ridge. If there is only one ridge this also means there are **no** cross-terms. But what about all the fragments or the undetectable situation of an auto-term and cross-term perfectly overlapping one another? Those do not matter. All of those issues boil down to the presence of more than one ridge. Being able to correctly reconnect the fragments of what should be a single ridge is beyond the scope, as those situations only occur if there are cross-terms. There can only be cross-terms present if the signal being tested is not a mono-component.

So ultimately, this problem appears to be nearly insurmountable without knowing how many components make up the signal being tested. But here again taking a different perspective and realizing that the only number of ridges that matter is one and more than one, we find ourselves with a practical, useable means for determining the mono-componentness of a signal!

3.5 Limitations and Usage

In the end, the actual algorithmic determination of mono-componentness results in a boolean decision. While, the ability to determine the exact number of components in a multi-component signal is not possible due to issues such as the undetectable situation of a cross-term occurring directly on top of an auto-term, still in the end the objective is to quantitatively determine if a signal is a mono-component and this can be done reliably because all that ultimately matters is if there is more than one ridge. Since a cross-term will only exist if there is more than one mono-component, the issue of a cross-term and auto-term coinciding is moot. But the measure is limited to reporting if a signal is a mono-component or a multi-component, where being a multi-component is determined by failing to be a mono-component.

The determination of mono-componentness does carry some limitations on its usage. It can only be applied to one-dimensional signals. Also, as with most things it is not appropriate for all situations. A specific example of this would be extracting source signatures from a hyperspectral image pixel, as the source spectrum of an “element” is not a mono-component signal. The metric is best suited for applications such as speech, communication signals, measurements (e.g., chemical concentration in the atmospheric, network port activation, particle detectors, etc.). A simple rule-of-thumb is that if the application is more acoustical in nature (i.e., waves) it is likely applicable. For example, the metric would not work for something such as mass-spectroscopy where the elemental returns are not conducive to being represented by a summation of mono-components, e.g. the spectral signature of an element cannot be composed of a sum of other elemental signatures.

The fact that this work created a free-standing measure of mono-componentness provides the maximum usefulness. An ideal usage of the mono-component measure is in combination with a single-channel signal decomposition method. By integrating mono-componentness into a single-channel decomposition algorithm the results can be quantitatively determined to be a mono-component or only partially decomposed.

In fact, the mono-component measure will see its own improvements through the reduced complexity resulting from even a poor quality decomposition. By decomposing the starting multi-component signal into a set of “potential” mono-components the number of ridges that the Crazy Climbers algorithm must find is lowered. Even if the decomposition results in the original signal only decomposed into two pieces the impact on the number of cross-terms contained in each of those two resulting signals is reduced from $\frac{N(N-1)}{2}$ to $\frac{((N-1)((N-1)-1))}{2}$, where N is the number of components. More simply put this results in $N - 1$ fewer cross-terms in the WVD for each of the decomposed signals. At any rate, the fewer ridges there are the faster the Crazy Climbers algorithm executes. An example of integrating the mono-componentness measure with a signal decomposition method is presented in Chapter 4.

3.6 Mono-component Measurement Case Studies

The mono-componentness measure was carefully designed such that the strengths of each key piece and its location in the algorithms processing flow allow the mitigation of the short-comings of the individual core elements. By applying the measure to *potential* mono-component candidates instead of the full multi-component signal, the TFR generated using the WVD contains only a few components at worst, thus shielding the interpretation and readability issue of the WVD when applied to multi-component signals. Next, by applying the Crazy Climbers algorithm to the TFR resulting from the WVD of the less complex potential mono-components, the number of potential ridges and the fracturing of ridges due to cross-terms is greatly reduced. This in turn aids the ridge detections by reducing the number of possible ridges (compared to the WVD of the original multi-component signal) and therefore simplifying the input for ridge counting. Because all of this occurs after passing through the proverbial gate keeper that is the WVD, we know that if only one ridge is present that the potential mono-component is in fact a mono-component.

Thus, the validity and functionality of the mono-componentness measure is based on a solid, verified foundation. Because of this only two examples of the use of the mono-componentness measure are shown: 1) Its response to a true mono-component signal. 2) Its response to a multi-component signal.

3.6.1 True Mono-component Case


In this case, the signal is single mono-component. It is a linear FM chirp of the form shown in Eq. (3.2)

$$y(t) = \cos(2\pi(\beta t^2 + f_0 t)) \quad (3.2)$$


and

$$\beta = \frac{f_{end} - f_0}{t_{end}}$$

where f_0 is the starting instantaneous frequency, f_{end} is the instantaneous frequency at time t_{end} . The true mono-component test signal was generated with $f_{end} = 0.3$ and $f_0 = 0.1$. A total of 1024 time steps were used making the effective sampling rate 2048 Hz. The TFR was generated using the highly efficient discrete WVD Matlab implementation of O'Toole [50, Chp. 6.6][34, 55]. The results of which are shown in Figure 3.6a. The intermediary results (occupancy measures) from applying the CC algorithm to the WVD TFR, Figure 3.6a, are displayed in Figure 3.6b, while the final ridge detection results generated by the occupancy chaining step of the CC algorithm is shown in Figure 3.6c. As is easily observed in all three plots in Figure 3.6 only a single ridge is present regardless of the point in the mono-componentness measure where the data is observed.



(a) The TFR of the true mono-component signal as produced via the WVD (b) The occupancy measure generated using Crazy Climbers on the TFR



(c) The ridges generated through the “chain-ing” of the occupancy values

Figure 3.6: The three key steps of the mono-componentness measure and its behavior when applied to a true mono-component. (a) The WVD of the true mono-component. (b) the occupancy measure produced by applying CC to the WVD in (a). (c) The ridge resulting from the final step of the CC algorithm which “chained” the occupancy values together to extract the ridge.


3.6.2 Multi-component Case

In this case, a multi-component signal is synthesized so that ground truth will be available. As with the mono-component test case, Section 3.6.1, the mono-componentness measure will be applied directly to the signal. The multi-component signal consists of three parallel linear FM chirps of the form shown in Eq. (3.2). Each of these chirps is a true mono-component and is generated with an IF that is a multiple of the first component's IF. This is done to create a multi-component signal that exhibits an especially difficult analysis situation when using the WVD. As mentioned previously in Section 2.6 a cross term occurs between each unique pairing of components. In this case, the cross term generated between the first and third component pair occurs directly on top of the location of the second component. This exact test signal was described in [56]. The specific beginning and ending IF of each mono-component is shown in Table 3.1.


Table 3.1: The starting IF, f_0 , and ending IF, f_{end} , values used to generate each of the three linear FM mono-components.

Mono-component	f_0	f_{end}
Comp #1	0.01 Hz	0.1 Hz
Comp #2	0.02 Hz	0.2 Hz
Comp #3	0.03 Hz	0.3 Hz

As with the mono-component test case, a total of 1024 time samples are generated for each component. The components are then summed to produce a three-component multi-component signal. The WVD is calculated of the multi-component signal as a whole and the results of which are shown in Figure 3.7a. Occupancy values from applying the CC algorithm to the TFR, shown in Figure 3.7a, are displayed in Figure 3.7b, while the final ridge detection results generated by the occupancy chaining step of the CC algorithm is shown in Figure 3.7c. Unlike for the true mono-component test case, Section 3.6.1, there are not simply three ridges found in the WVD of the multi-component test signal. In fact, we know the ground truth is that



(a) The TFR of a multi-component signal as produced via the WVD. (b) The occupancy measure generated using Crazy Climbers on the TFR in (a).



(c) The ridges generated by the “chaining” of the occupancy values in (b).

Figure 3.7: The three key steps of the mono-componentness measure and its behavior when applied to a true mono-component. (a) The WVD of the three component multi-component. (b) the occupancy measures produced by applying CC to the WVD in (a). (c) The ridges resulting produced by the CC algorithm once the occupancy values are chained to form extracted ridges.

there are only three ridges, but between the cross terms, the auto terms, and the overlapping cross and auto terms in the middle there are far more than just three ridge structures. CC found 11 ridges in this case. In regards, to the behavior of the mono-componentness measure when applied to a multi-component signal it has performed correctly and as intended. The multi-component signal failed the test of monosity.

Chapter 4

Applications and Analysis using the Mono-componentness Measure

The mono-componentness measure is a useful tool, but like a hose without a faucet or a nail without a hammer the mono-componentness measure needs signals to test. We need something to supply a set of “potential” mono-components to be evaluated using the mono-componentness measure. Given that one of the key ideas that resulted in the creation of the mono-componentness measure was to make it possible to leverage quadratic superposition to allow the construction of an interference free or nearly free TF representation, we naturally arrive at the need to decompose a multi-component signal. This combined with the other motivating objective of enabling the use of quadratic TF methods without the burden of selecting the “best” one blindly. That is, to eliminate or significantly decrease the prior knowledge of a signal needed in order to analyze the signal, and the scope of viable decompositions narrows to blind decomposition algorithms. In our case, the focus has been on decomposition algorithms that require only a single input signal.

The new ability to quantitatively evaluate the completeness of a decomposition revealed the startling degree of lacking in these methods. What had appeared to be adequate at least, was painfully exposed as being wholly unsuited to fulfilling the

desire to obtain a set of mono-components from a multi-component signal suitable for constructing an interference free WVD via superposition. It is clear that there is no blind, single-channel decomposition that comes close to decomposing a multi-component signal into an “adequate” number of mono-components, where adequate would ideally be a set of mono-components which containing the majority of the original signals energy. Because of this new and unexpected issue revealed by the first-of-its-kind mono-componentness measure, the focus changed from an interference free WVD to evaluating and improving decomposition algorithms.

While investigating decomposition algorithms the eventual determination was made that there are very few 1-D signal decomposition techniques available that do not require that the input data be reformatted into a matrix, through truncation and concatenation or through the resolution crippling application of STFT [57, 58]. The small number of algorithms that fit into the guiding motivation of requiring little or no prior knowledge of the signal or manipulation of the signal effectively limits the choices of blind 1-D decompositions to Empirical Mode Decomposition (EMD).

4.1 Single-Channel Decompositions

Algorithms and techniques for the decomposition of complicated signals are always of interest these days and have been for some time now. Two of the more popular techniques continue to be empirical mode decomposition (EMD) and independent component analysis (ICA). However, the ICA family of algorithms is not well suited for single-channel data sources due to the need of ICA to have at least as many channels (sensors) as there are independent sources [59]. There are a few published applications that attempt to modify/extend ICA to work with a single-channel time series, but in the end these extensions fall into at least one of the following:

1. Using a “pre-processing” step, usually STFT or spectrogram, to create a matrix to which ICA is then applied [60].

2. The single-channel signal is windowed or truncated to form a matrix [57].
3. Noise or some other processes are used to form additional, “virtual” channels from the single-channel signal [61, 62].

While there is nothing wrong with any of these approaches for applying ICA to single-channel/single-sensor data, they create the potential to produce misleading results in the case of non-stationary, transient data. The most obvious scenario for this is simply that the “window” used may segment the event that was captured. Second, is that our goal is to have high time and frequency resolution – something that the STFT and the spectrogram are not capable of. So if the potential components are extracted using STFT or the spectrogram the resolution has already been drastically limited and we can only lower the resolutions from that point on. However, Empirical Mode Decomposition (EMD) is a decomposition algorithm that is well suited for performing blind, single-channel decomposition task.

4.1.1 Empirical Mode Decomposition

EMD is a data driven decomposition method first proposed in 1998 by Norden Huang [63]. It is unique among decomposition methods due to its lack of need for any prior information about the signal being decomposed and its ability to operate on single channel data series without truncation, windowing, basis selection, or reformatting. This is particularly valuable when analyzing non-stationary, multi-component signals. Signals for which little to nothing is known regarding its contents, except the assumption that it contains an event of interest – single-shot signals.

EMD produces a set of what are called intrinsic mode functions (IMFs) which when summed exactly equal the original input signal, where an IMF is defined as a signal which meets the two requirements: 1) The total number of zero crossings and the total number of extrema points must be equal or differ at most by one, 2) The mean of the maxima envelope and the minima envelope must be zero everywhere [63].

The EMD algorithm operates iteratively using two loops. The outer loop results in the decomposition of the input signal into a set of IMFs. The inner loop successively refines the current IMF “candidate” that is evaluated against the IMF criteria, embodied as a set of stopping criteria. When the outer loop completes an iteration is when the inner loop produces a candidate that meets the stopping criteria.

The EMD inner loop fits one spline between the local minima, $e_{min}(t)$, and another between the local maxima, $e_{max}(t)$, of the input signal $x(t)$. The average envelope, $m(t)$, of the two splines is calculated and subtracted from the signal that $e_{max}(t)$ and $e_{min}(t)$ were derived from, resulting in a new signal, $d(t)$. This process repeats, called “sifting,” until the stopping criteria are met by the current $d(t)$ signal. At this point the signal remaining, $d(t)$, meets the criteria to be an IMF (with some variance allowed by the stopping criteria).

The stopping criteria determine how close to these criteria the candidate IMF actually is before beginning the next IMF sifting cycle. The next cycle of sifting begins with the original, starting signal of the previous sifting round minus the IMF just obtained. So in the case of transitioning from the first IMF sifting cycle to the second it would be $x_1(t) = x_0(t) - IMF_1$, where $x_0(t)$ is the original signal, IMF_1 is the IMF resulting from the first sifting cycle, and $x_1(t)$ is the starting signal for the second IMF sifting cycle.

Algorithm:

1. Identify all extrema of $x(t)$
2. Interpolate between minima (respectively maxima), ending up with some envelope $e_{min}(t)$ (respectively $e_{max}(t)$)
3. Compute the mean, $m(t) = \frac{(e_{max}(t) + e_{min}(t))}{2}$
4. Subtract mean envelope to produce the next iterations input signal $d(t) = x(t) - m(t)$

Steps 1-4 form the inner loop and “sift” the input signal until $d(t)$ is zero-mean (or within the user set tolerance of it)

5. If the stopping criteria are met then $d(t)$ is a “mode” (IMF)
6. The IMF cycle is repeated using the residual $r(t) = x(t) - d(t)$, as the “new” $x(t)$, such that each successive IMF sifting cycle’s starting signal is the original $x_{initial}(t)$ minus all previously extracted IMF signals.

This cycle continues where each successive sifting process operates upon the original input signal minus all IMFs extracted so far. The sifting process is halted when the initial $x(t)$ signal minus all the extracted IMFs contains only three extrema and/or meets the IMF criteria without the need for even a single sifting cycle. The stopping criteria consist of three parameters defined to allow some tolerance with respect to the zero mean envelope IMF criteria. This requirement is especially difficult to meet perfectly due to the iterative spline fittings. To account for this, along with the fact that continual pursuit of perfect convergence produces less and less benefit while consuming more time and resources. As such, in [64] the three parameter tolerance based stopping criteria were purposed and quickly became the norm among the users of EMD. These “classic” stopping criteria are $\theta_1 = 0.05$, the global fluctuation allowance, $\theta_2 = 0.5$, the local fluctuation allowance, and $\alpha = 0.05$, where $(1 - \alpha)$ determines what is “global” and α what is local in terms of the signal’s duration.

There are two alternate stopping criteria that can be used in addition to, or in lieu of, simply meeting the two IMF criteria of having a zero mean envelope and that the total number of extrema equals or differs by one the total number of zero crossings. These are: 1) use a fixed number of inner loop iterations, 2) extract a specified number of IMFs, i.e., a fixed number of outer loop iterations. These additional criteria are useful when some prior information is available to suggest how many modes (components) may be present in the multi-component signal. Also, as EMD operates empirically using only information derived from the extrema and zero crossings of a signal, combined with spline fitting, there is an inherent degree of amplitude error in the IMFs it produces. There are always IMFs extracted that

will never pass the mono-component measure and have no physical meaning, but these IMFs can be distinguished through simply measuring the percentage of the original multi-component signal's energy each IMF contains and segregating out the IMFs that fail to pass some minimal energy content threshold. These “junk” IMFs are the result of the accumulation of small amplitude errors as IMF sifting cycles progress. These amplitude errors are due to the “masking” of extrema points due to a weaker amplitude component's extrema being dominated by a component with a higher amplitude, even if this dominate amplitude only exist for a short duration [65, 66, 67].

EMD is very attractive due to its dynamic data driven nature. It requires no windowing or pre-processing sort of data manipulation. However, it does have several issues.

4.1.2 Known Limitations of EMD

One of the most often mentioned negatives of EMD is that it does not have strong mathematical foundation, which is true [64]. However, the fact that EMD does not depend on *a priori* knowledge of the signal and is capable of operating on non-stationary signals outweighs the negatives regarding its lack of a rigid mathematical base because of its algorithmic style of operation. There are two issues with EMD which do have a meaningful impact on its results that cannot be overlooked. They are: 1) mode-mixing and 2) the separability limit. The first issue has been studied and several solutions have been proposed to greatly limit its occurrence. The second is not necessarily an issue so much as a fundamental characteristic of the EMD process regardless of the particular variant of EMD that is used and thus is also an inherent limit on the work present in this document.

Mode-Mixing

Mode-mixing is the name given to an issue that can be seen in the IMFs resulting from the use of standard EMD. It describes the situation where signal content from one IMF is clearly present in another, almost as if a piece from each of the mode mixed IMFs was transcribed between the two. Mode-mixing is very obvious when examined visually as seen in Figure 4.1.




Figure 4.1: An example of mode-mixing in the potential mono-components (IMFs) resulting from the application of standard EMD. (a) The three components summed to create the multi-component signal that EMD was used on. (b) The six potential mono-components resulting from EMD. Notice the obvious mode-mixing between IMF-1 and IMF-2, and IMF-2 and IMF-3.

Mode-mixing stems from the ambiguity involved in the EMD sifting process, in particular, the extrema used to form the maxima and minima envelopes. Since, EMD is data-driven there is not an algorithmic way to determine whether a particular maxima or minima belong to a particular mode/frequency, especially if the mixed modes are harmonics of each other and/or phase synchronous. Effectively, what is occurring when there is mode-mixing is the same as aliasing in discrete sampled signals. Aliasing and mode-mixing have the same end-effect, but occur through

entirely different mechanisms. That is to say, unlike aliasing, sampling at a higher rate has no impact on the prevention of mode-mixing.

If in the signal being decomposed there exist local maxima and minima, whether a local extrema came from a single component or the summation of any number of components is unknown. When two components have extrema in close proximity their summation results in a single extrema point. When EMD sifts out that “merged” extrema it is no longer available for reconstructing the remaining component. Thus, the mode-mixing is particularly likely to occur when the signal being decomposed contains phase-synchronous harmonics like shown in Figure 4.1 where the 1Hz and 2Hz components have resulted in mode-mixing across the first three IMFs.

Separability of Closely Spaced Frequencies

EMD is quite powerful given that it requires no prior information, but as one should expect it has a limitation on what it can decompose. Specifically, on how closely spaced two tones may be and still be separable. In fact, a fantastic simulation study was conducted by Rilling and Flandrin in [65] that is still used to evaluate the ability of a decomposition algorithm to separate two closely spaced tones [68, 69, 70]. The groundbreaking work by Rilling and Flandrin systematically explored the effects that frequency differences and amplitude differences have on EMD’s ability to separate two tones.

The framework developed makes use of the fact that it is the ratio between the two signals’ amplitudes and frequencies that impacts EMD’s ability to resolve the signal as two frequencies or a single frequency. Starting with Eq. (4.1) and recognizing that the formulation can be simplified to aid calculation without loss of generality by using Eq. (4.4) to represent a two-tone signal using two parameters: a the amplitude ratio and f the frequency ratio. Note that it is assumed, and easily ensured since it is a simulation, that the sampling frequency is much greater than either f_1 or f_2 . Along this same line, one of the tones is chosen to be the higher frequency tone and fixed at 1Hz and with a peak-to-peak amplitude of two (+1 to -1). Because it is an arbitrary

choice, a_1 and f_1 will belong to the high frequency signal with the fixed frequency and amplitude allowing the amplitude and frequency ratios to be defined by Eq. (4.2) and Eq. (4.3), respectively. This also lets the span of f only need to be swept from $(0, 1)$.

$$x(t) = a_1 \cos(2\pi f_1 t) + a_2 \cos(2\pi f_2 t) \quad (4.1)$$

$$a \triangleq \frac{a_2}{a_1} \quad (4.2)$$

$$f \triangleq \frac{f_2}{f_1} \quad (4.3)$$

$$x_{(a,f)}(t) = \cos(2\pi t) + a \cos(2\pi f t) \quad (4.4)$$

As EMD extracts higher frequencies (faster oscillations) first, the separability metric compares the difference between the first IMF extracted and the original, higher frequency, component normalized by the other original component as shown in Eq. (4.5).

$$m(a, f) = \frac{\| \text{IMF}_1(t) - \cos(2\pi t) \|}{\| a \cos(2\pi f t) \|} \quad (4.5)$$

To illustrate the separability limit, we have reproduced the same experiment as Rilling and Flandrin using a fixed 10-iterations per IMF sifting. Figures 4.2 and 4.3 show the separability metric scores when using a fixed 10-iterations and 100-iterations respectively, where zero is a perfect separation and one indicates that the two tones could not be separated. Figures 4.4 and 4.5 show the same results as Figures 4.2 and 4.3 but from an informative 3-D perspective. Finally, in Figures 4.6 and 4.7, we see the same two tone separability test using the “classic” EMD mean envelope tolerance stopping criteria suggested in [64]. This particular test has not previously been published until now. As you can see it has particularly terrible separation characteristics even compared to simply performing ten brute force iterations regardless.

Ingrid Daubechies said it well in [5] that, “In summary, the EMD algorithm has




Figure 4.2: The two-tone separability score for EMD when using a fixed 10-iterations per IMF sifting. The closer to zero the better the two-tones where separated, while one and higher is a failed separation.




Figure 4.3: The two-tone separability score for EMD when using a fixed 100-iterations per IMF sifting. The closer to zero the better the two-tones where separated, while one and higher is a failed separation.




Figure 4.4: The two-tone separability score for EMD when using a fixed 10-iterations per IMF sifting viewed in 3-D.




Figure 4.5: The two-tone separability score for EMD when using a fixed 100-iterations per IMF sifting viewed in 3-D.




Figure 4.6: The two-tone separability score for EMD when using the “classic” tolerance based stopping criteria, $[\theta_1 = 0.05, \theta_2 = 0.5, \alpha = 0.05]$.




Figure 4.7: The two-tone separability score for EMD when using the “classic” tolerance based stopping criteria viewed in 3-D.

shown its usefulness in various applications, yet our mathematical understanding of it is still very sketchy.” Despite the issue of mode-mixing and the limitation on the separability of tones, EMD is a valuable tool for the decomposition of non-stationary, multi-component signals. Its methodology does not require any prior information about the signal to be decomposed and it does not perform any operations that would pre-maturely limit the resolution of decomposition results. However, a serious issue with EMD (and any single-channel decomposition algorithm) is that the resulting components have no guarantees about their mono-componentness. The decomposition results may be orthogonal or statistically independent or even possibly contain mono-components, but there is no assurance they any will be mono-components.

While the objective, algorithmic means for determining if a signal is a legitimate mono-component, or not, is in itself a new and unique creation its greatest benefits are in its addition to existing applications. The most powerful and beneficial is its addition to EMD. EMD and its related formulations, Ensemble EMD (EEMD) [66], Complete EEMD with additive noise (CEEMDAN) [71, 72], the optimization based EMDs [70, 69, 73], etc. are all powerful and useful tools. Especially, when there is little prior information, unique one-of-kind signals, etc. but regardless of its formulation none of these return IMFs with any guarantees beyond that the IMFs have a mean of zero and at most a difference of one between the number of extrema and the number of zero crossing. Depending on the particular EMD formulation determines how strictly it adheres to those two requirements or some constraint encapsulating those objectives in the case of optimization based EMD algorithms. In the end, there is no single-channel decomposition method that produces a set of only valid mono-components. Not because we don’t want mono-components, but because we have had no way of knowing for certain if we had mono-components. By combining the mono-component measure with EMD a new decomposition method with quantifiable results was created.

4.2 Empirical Multi-Component Decomposition

While a number of publications refer to EMD or one of its variants as producing mono-components [67, 74, 75], and EMD is capable of this, but not because its algorithm and design require or enforce that it does. IMFs are fundamentally defined by two requirements neither of which require that the IMF be a mono-component in order to meet them. In fact, IMFs are typically not mono-components [76]! By adding the mono-componentness measure, IMFs that pass the measure are now endowed with physical meaning beyond that of just having a mean envelope of zero and that the number of zero-crossings and number of extrema differing by one or less. They now are verifiably, actual mono-components!

Incorporating the mono-componentness measure into EMD as part of the stopping criteria also opens a new vector of improvement in EMD. The original implementation allows a higher fixed number of inner loop iterations to be used, but now with the ability to escape out of the inner loop if the IMF candidate passes the mono-componentness test instead of executing the full number of iterations for every IMF, it allows the advantages of improved separability without the risk of over sifting the IMFs.

Unfortunately, EMD is not without its own issues as mentioned in Section 4.1.1. However, the direct incorporation of the mono-componentness measure into EMD-type algorithms as a stopping criteria allows us some advantages. The issue of the frequency separability can be addressed to some degree by allowing the number of inner sifting iterations to be increased, a known means for raising the limit of separation to be above the 2/3 frequency ratio as documented in [65], see Figure 4.3 versus Figure 4.2. By incorporating mono-componentness as a stopping criteria the maximum inner sifting iteration can be set relatively high, but now not all of those iterations must be executed. For if at any time the signal passes the mono-componentness measure the inner sifting operation will be terminated, returning an IMF, a true mono-component IMF no less. The mono-component measure allows the

automatic escaping of an IMF sifting cycle, while retaining the desirable characteristic of allowing the separation of some frequencies that have a frequency ratio greater than two-thirds.

For there is a distinct danger of executing too many sifting iterations when extracting an IMF. The issue is that the resulting IMF becomes a signal with a constant amplitude and a constant frequency modulation, removing any amplitude variations along with any physical meaning [77]. To this end a new EMD-based blind single-channel decomposition is created by integrating the mono-componentness measure into what we will call Empirical Multi-Component Decomposition (EMCD). This provides several advantages:

1. Higher number of inner iterations for refining IMFs, but only the number necessary to sift a mono-component will be performed.
2. The passing IMFs are without a doubt mono-component signals — a truly meaningful physical attribute.

4.2.1 An Advantageous Combination

Each core piece that makes up the unique mono-component measure presented in this work has its own shortcomings. The WVD has a terrible and pervasive interference problem that comes from the very core of its formulation. EMD blindly decomposes a signal into a set of IMFs of which the criteria met to qualify as an IMF is only useful for declaring a signal to be an IMF and nothing else. The Crazy Climbers algorithm applied directly to the WVD results of a multi-component can easily extract an unwieldy multitude of ridges from a TFR, and in the presence of interference terms most ridges are insignificantly short and/or result in a ridge that incorrectly chains a large number of tiny ridges into a meaningless path.

Now, if the signal source to the mono-component measure is a set of IMFs, even if the IMF is actually another multi-component signal, it has at the worst at least one less component-like signal in it and more typically (though not guaranteed)

significantly fewer component-like signals than the original multi-component input to EMD. Based on the extrema fitting and sifting of EMD, a spline will be fitted and removed so long as the stopping criteria are not met. The very separation evaluation metric of [65] uses this behavior to determine the separability of two tones. By taking each IMF, which might be a legitimate mono-component but at the worst is a multi-component signal of fewer components, and calculating the TFR using the WVD, the number of interference terms produced by each IMF will be less than the number of interference terms produced by the multi-component signal as a whole. Taking the TFR and applying Crazy Climbers likewise results in fewer spurious ridges existing and improves the ridge detection. Each step of the algorithm reduces the complexity of the input to the next step.

It is through the careful design and development of the mono-componentness measure that allows the strengths of each core piece to mitigate the weakness of the other core pieces and result in a capability never before available and with its incorporation into EMCD both algorithms benefit. EMCD is capable of producing true mono-components, while the mono-component measure benefits from the reduced number of components in each input even if they are not valid mono-components.

4.2.2 An Analysis of the Mono-componentness of Empirical Mode Decomposition's Intrinsic Mode Functions

The use of synthetic test signals has long been an accepted means of demonstrating the capabilities and issues of a TF distribution and as such several of these multi-component test signals will be used to illustrate the power and usefulness of the presented mono-component measure. For this detailed examination of the power of the mono-component measure as both a quantitative check of EMD's results and in the form of a new EMD variant in which mono-componentness is part of the stopping criteria, we will make use of a particularly problematic synthetic test signal.

The signal used is carefully constructed to not only create an unfilterable

interference problem due to the direct and complete overlap of an auto-term and a cross-term, but also uses signals that are on the wrong side of EMD’s separation limit, see Section 4.1.2. The specific signal to be used was described in detail in Section 3.6.2.

First, the separability of EMD using the traditional stopping criteria is demonstrated on this signal. The criteria consists of three parameters defining the tolerance allowed when it comes to an IMF having a zero mean envelope as first proposed in [64] and quickly adopted by the users of EMD. As stated in Section 4.1.1, this “classic” stopping criteria are $\theta_1 = 0.05$, the global fluctuation allowance, $\theta_2 = 0.5$, the local fluctuation allowance, and $\alpha = 0.05$, where $(1 - \alpha)$ determines what is “global” and α what is local in terms of the signal’s duration.

The 3-component signal is shown in Figure 4.8. To most clearly illustrate the results of the EMD decomposition we show first the time domain plots of the input signal, each IMF extracted, and finally the residue. This is shown in Figure 4.9. Also, the WVD of the 3-component signal as well as the WVDs of each individual mono-component used to construct the 3-component signal are shown in Figure 4.2.2.

Immediately, followed by the results of testing the IMFs and residue using the mono-componentness measure, as shown in Figure 4.11. These results are displayed using two of the intermediary representations used internal to the mono-componentness measure: the WVD and the ridge detection map. The ridge detection map contains an informative label providing the number of ridges found and the number of iterations EMD ran to produce the IMF.




Figure 4.8: The time domain plot of the 3-component test signal.




Figure 4.9: The time domain plot of the input to EMD (topmost pane), followed in descending order by the first IMF extracted ending with the residue (bottommost pane).




Figure 4.10: The WVD of the 3-component multi-component signal along with the WVD of each individual mono-component used to form the multi-component signal, i.e., the ground truth mono-components.




Figure 4.11: The intermediary calculations created internally by the mono-component measure when the IMFs of the 3-component test signal was decomposed using the standard EMD stopping tolerances which involves no maximum iteration limit. The WVD of the original 3-component signal, all the IMFs, and the residue are displayed with their associated ridge detection map labeled with the number of ridges detected and the number of iterations EMD performed to extract that particular signal below each WVD.

Moving right along to the next test, we switch from using the mean envelope tolerance stopping criteria to using a fixed number of sifting iterations (inner loop). The original separation test by Rilling and Flandrin used a fixed 10 iterations. We will begin with this but also then repeat the experiment using 100 iterations. As we saw before in Figures 4.2 and 4.3 the higher the number of sifting iterations the better EMD does at separating closely spaced tones; at the expense of executing more iterations of course. A consequence of using a larger number of sifting iterations is that as the number of iterations increases the closer the extracted IMFs will be to becoming constant amplitude signals, which results in the loss of the meaningful amplitude information [77].

The time domain plot of the IMFs generated using a fixed 10 iterations is shown in Figure 4.12 along with the revealing intermediary steps extracted from the mono-component measure in Figure 4.13 showing the WVD, ridge detection map, ridge counts and iterations executed for each IMF. Likewise, the corresponding results for EMD using a fixed 100 iterations are shown in Figure 4.14 and Figure 4.15. Where things really start to get interesting is in Figure 4.16 which shows the time domain plot of the IMFs extracted using EMCD with a maximum allowed number of iterations of 100, but with the addition of the mono-componentness measure to allow the inner sifting to stop as soon as a valid mono-component is sifted. The WVD, ridge map, ridge counts, and number of iterations executed for each IMF returned by EMCD are shown Figure 4.17.




Figure 4.12: The time domain plot of the input to EMD (topmost pane), followed in descending order by the first IMF extracted and ending with the residue (bottommost pane), when EMD's stopping criteria is to sift a fixed 10-iterations each IMF cycle.




Figure 4.13: The intermediary calculations created internally by the mono-component measure when the IMFs of the 3-component test signal was decomposed using a fixed 10-iterations per IMF. The WVD of the original 3-component signal, all the IMFs, and the residue are displayed with their associated ridge detection map labeled with the number of ridges detected and the number of iterations EMD performed to extract that particular signal below each WVD.




Figure 4.14: The time domain plot of the input to EMD (topmost pane), followed in descending order by the first IMF extracted and ending with the residue (bottommost pane), when EMD's stopping criteria is to sift 100-iterations each IMF cycle.




Figure 4.15: The intermediary calculations created internally by the mono-component measure when the IMFs of the 3-component test signal was decomposed using a fixed 100-iterations per IMF. The WVD of the original 3-component signal, all the IMFs, and the residue are displayed with their associated ridge detection map labeled with the number of ridges detected and the number of iterations EMD performed to extract that particular signal below each WVD.




Figure 4.16: The time domain plot of the input to EMD (topmost pane), followed in descending order by the first IMF extracted and ending with the residue (bottommost pane) when using the new EMCD that incorporates the mono-componentness measure and the maximum number of sifting iterations set to 100.




Figure 4.17: The intermediary calculations created internally by the mono-component measure when the IMFs of the 3-component test signal was decomposed using the new EMCD algorithm with the maximum number of sifting iterations set to 100. The WVD of the original 3-component signal, all the IMFs, and the residue are displayed with their associated ridge detection map labeled with the number of ridges detected and the number of iterations EMD performed to extract that particular signal below each WVD.

As you can see in Figure 4.17 for the IMFs that are valid mono-components the number of iterations needed is far less than the 100 iterations needed to extract the same number of valid mono-components as the standard EMD with a fixed 100 iterations per IMF.

4.2.3 Problematic Cases

There are a handful of multi-component signals that are known to be problematic for both decomposition and/or ridge detection. These weaknesses are *not* due to the mono-componentness measure, but primarily an issue in EMD. These problematic cases can be grouped into the following general categories:

1. Tones/frequencies whose frequency ratio is, $\frac{f_1}{f_2} \geq \frac{2}{3}$ and/or whose amplitudes are highly disproportionate, such as $\frac{a_1}{a_2} = \frac{0.1}{1}$.
2. Crossing IFs. In particular, a linear FM chirp whose IF is increasing with time and an identical linear FM chirp whose IF is decreasing with time will form an “X”. Since, the signals are identical with respect to the rate of frequency change the IMFs extracted via EMD are a “V” shape instead of a diagonal line. Although, a correct extraction of sorts can be obtained, as all the frequency content and change is represented, the time of occurrence is only correct for half the extracted component.

4.3 Quantitative Measure of the Quality for the Hilbert-Huang Transform Spectrum

As the research and investigations that eventually gave birth to the mono-componentness measure progressed, several other quite novel developments also emerged. One unique and useful contribution of particular note that results directly from the ability to

determine mono-componentness is the new ability to determine the accuracy and quality of a signal’s Hilbert-Huang Transform (HHT) spectrum [63, 78].

EMD was developed in order to create near-mono-component signals from the input data for the HHT. The HHT spectrum is a time-frequency result that is intended to display the instantaneous frequencies (IFs) of each component making up a multi-component signal. The components are actually IMFs which may or may not be legitimate mono-components. The IF of each IMF is calculated via the Hilbert transform. The accuracy of the resulting IF is inescapably bound to the mono-componentness of the signal supplied to the HHT [78].

A dramatic illustration of the importance of using valid mono-components can be seen in the HHT spectrum. The left plot in Figure 4.18 shows the HHT spectrum calculated using the IMFs, while on the right is the HHT spectrum calculated using the ground truth mono-components; the difference is striking. The bottom IF ridge is correct in both plots as in both cases since it is a valid mono-component. The middle and top IF ridges in the left HHT spectrum are corrupted due to the incomplete separation of the two components as can be seen in the WVD of IMF-1 in Figure 4.11. While IMF-2 is a mono-component, an unknown portion of its amplitude was left behind in IMF-1. Regardless, the reliability of an IF estimated using the Hilbert transform quite clearly requires the use of a legitimate mono-component for the best results.

A simple, yet robust, measure of the quality and accuracy of the content of the HHT spectrum can be calculated by summing the number of IMFs found to be valid mono-components divided by the total number of IMFs extracted, as in Eq. (4.6). Where the closer to one the resulting ratio is the more reliable and/or accurate the IFs represented in the HHT spectrum will be. For example, the ground truth HHT spectrum shown in Figure 4.18b is 1 (as expected), while the score for Figure 4.18a is 0.28.




Figure 4.18: HHT spectrum of (a) IMFs generated using EMD on the 3-component multi-component signal, whose quality metric is 0.28, (b) the three mono-components used to generate the multi-component, i.e., the HHT spectrum of the ground truth components, whose quality metric is 1.

$$HHTmetric = \frac{\sum ValidMono - components}{(\#ofIMFs + Residue)}, HHTmetric \in [0, 1] \quad (4.6)$$

As a closing thought, one could consider Eq. (4.6) to be biased, as the residue returned by EMD can never be a mono-component since by its very construction EMD terminates when the remaining signal no longer contains the minimum required three extrema necessary to have an oscillation at even the slowest of periods. Therefore, one may wish to only divide by the total number of IMFs extracted, discarding the residue when calculating the HHT metric. The results will still be the inclusive range of zero to one, where one is perfect.

Chapter 5

Discussion

5.1 In Summary

There are quadratic TF distributions that make strides toward suppressing their cross-terms, but with a high cost in prior knowledge of the signal being analyzed. We have focused on the WVD because it is the base of all quadratic TF distributions and also does not require any selections or parameter adjustments *a priori*, but it still suffers from cross-terms.

Ultimately, this work resulted in taking a unique approach to the cross-term problem. By focusing on mono-components instead of multi-components and seeking an alternative to mathematically eliminating the cross-terms. This resulted in seeking to use quadratic superposition to construct a high resolution, low interference TFR. In order to use superposition though one must have a set of mono-components which sum to equal the multi-component signal. Thus, the importance of the mono-component and the lack of a mathematical representation of this concept came to bare.

The final result is the creation of the first-of-its-kind quantitative measure of mono-componentness. By enabling the validation of a set of potential mono-components as legitimate or not the door is now open to allow the use of quadratic superposition to construct simultaneous high resolution TFRs of multi-component signals. Also, the

ability to objectively evaluate blind decomposition algorithms was realized as well. Unfortunately, this new ability disproved the assumptions regarding the EMD's ability to produce mono-components from real-world signals. Still, the door to creating a brand new signal channel blind decomposition algorithm that takes into consideration mono-componentness is now wide open.

Also, by eliminating the need to visually examine TFRs to determine mono-componentness or even the quality of the HHT TFR is no longer needed allowing much longer series to be processed in the TF domain. Time-series as long as 10,000 samples were regularly processed as part of this work. Given greater computational resources (multi-core CPUs, graphic processing units, etc.) even longer series can be processed and evaluated without requiring the even more computational expensive step of rendering the data as an image. A highly subjective and error prone task just from the simple perspective of screen resolution.

Ultimately, the goal of a QTFD which requires no prior information was not achieved due to the new ability to quantify the results from blind decomposition algorithms. However, at the same time a path for their improvement was also created, all without violating the objective of requiring no prior information.

The mono-componentness measure requires no prior information and provides a quantitative, objective, and flexible means for detecting and testing for a key signal characteristic **that has no mathematical representation**.

This first-of-its-kind measure allowed the following additional first to occur:

1. A quantitative, objective method for evaluating the quality and accuracy of Hilbert-Huang Transform spectra.
2. A means for determining *if* the IMFs produced by EMD (or its variants) are actually mono-components, which has been typically only been assumed.
3. A viable constraint that will allow a decomposition algorithm to have an objective function representing mono-componentness.

4. Mode-mixing in EMD can be detected algorithmically since a mode-mixed IMF can never be a mono-component.

5.2 Future Work

There are several avenues for future work. The extension of the current mono-componentness measure to handle inner interference being the most important. The top ideas for future research and development are:

- Detection of inner interference using Hough transform to find concavities whose foci are all aligned.
- Possibility that higher-order WVD (e.g., 4th order, 6th order) might ideally localize mono-components with polynomial phase (i.e., mono-components that produce inner interference).
- Possibility that the Atles Q-distribution (member of the hyperbolic quadratic TF family) does not produce inner interference terms for mono-components with concavities in its IF law.
- Automatic detection of tones too closely spaced for EMD to separate.
- Investigate the effect noise has on the use of the measurement as EMD is known to struggle when SNR is poor.
- Improved ridge detector.
- Improved computational speed.

Some of these ideas are more mature than others and are discussed, along with any initial investigations which have already been performed, in the following sections.

5.2.1 Algorithmic Inner Interference Detection

There is the issue of inner interference that still needs to be addressed. The current method of ridge detection and counting to determine mono-componentness will fail in the presence of a mono-component, if that mono-components instantaneous frequency (IF) contains concavities [24], i.e., a mono-component with a polynomial phase function that is of an order greater than two. There are several straight forward options that can extend the current mono-componentness measure to deal with inner interference.

One is by detecting the multiple nested concavities of a mono-component with inner interference. By leveraging the fact that the inner interference terms generated in a concavity share the same concave shape as their source and spoon together the concaves can be detected, the focus of each calculated, and if they form a linear path then they are likely inner interference terms. The first step of which is to analyze the ridges already detected for curvatures such as with the Hough transform configured for parabola detection. If multiple curves are found the focus of each will then be calculated. If the foci are along a linear path this indicates that these curves are “nested” together and are likely inner interference terms. To ensure that the inner interference is the source of the multiple ridges all that is necessary is to remove the total number of ridges that were curved and whose foci fall along the same linear path from the total number of ridges detected. If the reduction in the ridge count results equals zero then it is inner interference in a mono-component that caused the TFR to contain multiple ridges, otherwise it a multi-component signal.

A second possibility would simply require that the “measure” and ridge counting be done twice. One pass as presented already, using the WVD, and once using a hyperbolic QTFD, such as the Altes-Marinovich Q-distribution [79] and performing ridge detection on its resulting TFR looking for a single ridge.

Similar to the previous possibility, is the use of what are known as the L-class distributions which are a generalized extension to Cohen’s class. The L-class

distributions are a higher order distribution (i.e., the order greater than quadratic) and display significantly less inner interference for signals having higher order (non-linear) IF laws [4]. The L-class distributions however have a much more serious outer interference issue though. This is why the combination of an L-class distribution and the use of the presented mono-component measure are an ideal environment for detecting inner interference. By integrating an L-class distribution into the mono-component measure as a second stage test for signals with more than one ridge, the complexity of detecting inner interference using the WVD can be simplified to detecting the presence of a single ridge in the L-class's TFR in the case where the WVD's TFR produced multiple ridges (all internal to the mono-componentness measure).

In fact, there is a possibility that thanks to the elegantly simple approach used in the design of the mono-componentness measure the WVD could be simply replaced with a fourth or sixth order L-class distribution entirely even though this family of distributions suffers worse outer interference issues when analyzing multi-component signals, a mono-component does not produce outer interference.

5.2.2 Improved or New Decomposition Method

One of the greatest challenges that emerge once the mono-component measure had been created and implemented was the fact that the measure was created to allow the use of the WVD's superposition property to enable an outer interference free WVD of multi-component signals to be created that also eliminate the need for any prior knowledge of the signal or the subjective selection of the "best" QTFD. However, once the mono-componentness of signals could be quantified it was suddenly very clear that the one-dimensional decomposition algorithms available today seldom produce any mono-components at all. Then add the additional complicating criteria of avoiding any decomposition method that requires prior knowledge or multiple instances of the signal and all that remained was EMD.

However, finally having an quantitative, objective means for determining if a signal is a mono-component has become somewhat of a double-edged sword. After all, it is the result of a desire to enable the practical usability of TF analysis to meet real-world challenges. Using the new ability to verify what visually appeared to be IMFs that were mono-component extracted from real-world signals it becomes shockingly clear that EMD fails to extract any true mono-components the majority of the time from real-world signals. A vast variety of real-world signals were decomposed and then tested using the developed mono-component measure and none of them produced a single mono-component. The signal tested where:

1. Bluetooth RF recording.
2. Bat echo-location chirp.
3. Ultrasonic pulse from a non-destructive ultrasonic concrete inspection tool through a known perfect concrete test specimen.
4. Voltage signal recorded from the micro-grid side of the common coupling point.
5. A single note played on a piano.
6. A dolphin echo-location ping.

Ultimately, they all suffered from either having components with drastically disproportionate amplitudes or frequencies whose ratio placed them squarely on the wrong side of EMD's separation limit. In fact, this resulted in the first real-world signal the author has ever seen that was un-separable not because the frequencies where too close but because the amplitudes where so unequal, despite the frequency ratio being well below the $\frac{2}{3}$ rds threshold at an ideal ratio of only $\frac{1}{3}$ rds.

Perhaps a structural health monitoring vibration signal would have had success, but the overall evidence is that there is still much work to be done in the realm of blind one-dimensional decomposition methods and the ability to quantify the completeness of a decomposition provides insights that have never before been available.

5.2.3 Computational Complexity

The mono-component measure is computationally expensive at this time. It is the calculation of the WVD that is the bottleneck. For short duration series, i.e., series of a couple thousand samples, the wait time is relatively minor. However, when processing time series containing more than a few thousand samples it can be quickly become frustratingly slow especially when the measure is integrated into an EMD variant as a stopping criteria. An implementation based on the superior QTFD implementation by O’Toole [31, 32, 37, 38] was created that made use of a high performance graphics processing unit to calculate the WVD of a large time series. Large is a relative term. In the time-frequency domain most published analysis use short signals by anyone’s definition. In fact, [80] make a point of how their method opens the door to “big data” in that they were able to calculate a TFR of a 512 sample signal in approximately 4 seconds. The graphics processing implementation of the WVD can calculate the full resolution TFR, full resolution being an FFT length equal to the signal length, of a 8384 sample signal in 2.9 seconds. Still, large is relative in this case, but one key need for processing large sequences is that they be rendered visually as little as possible, preferably only when something meaningful is automatically detected.

5.2.4 Parameter Values

The mono-componentness measure has a relatively small number of internal parameters that affect its sensitivity with regards to when a signal is deemed a mono-component signal. The majority of these control the Crazy Climbers ridge detection algorithm and include parameters such as the minimum length a ridge must extend, the number of “climbers” to use, and a minimum occupancy value.

The minimum occupancy value threshold is of obvious importance, but the effect of more minor parameters such as the step size of the crazy climber paths versus the step size used in the chaining process may be worth sweeping to provide a rule-of-thumb

for their setting. While setting the initial temperature in simulated annealing is an open-question and since CC is at its core a modified simulated annealing framework the question of whether to use a constant temperature versus a cooling schedule may be worth investigating. A quick contemplation of using a fixed temperature seems as if it might be useful, as it would allow a climber full mobility throughout its entire climb and since CC is designed to travel along/seek local maxima it may be beneficial to allow the climbers to travel more broadly helping to avoid the tendency for climbers to concentrate at the global maximum.

However, while running the two-tone separation test from Rilling and Flandrin [65] is the ideal test scenario for performing these sort of parametric sweeps they are very computationally costly to execute. A single two-tone signal from the test that has poor separability when only a single sifting iteration is performed, but separates well when 10 sifting iterations are performed was taken from the larger test frame and the minimum occupancy threshold was swept from 0.3 to 0.01 with a step of 0.01 while increasing the number of sifting iterations from 1 to 10 on a 32-core 2.3GHz computer took 17 minutes. This doesn't sound too bad until you scale to the full framework which has 8,100 two-tone signals and would ultimately take 95 days of continuous run time on the same system.

5.2.5 A Better Ridge Detector

While the Crazy Climbers algorithm provides the functionality required in order to implement the mono-component measure and does have a pedigree to back up its approach it does not mean it is the best or only option. In fact, in the software suite [81] provided by Carmona to accompany his book [52] contains several additional ridge detector algorithms that are completely undocumented or published other than their inclusion as source code in the Swave software package.

Despite the detailed information provided for the CC algorithm in [52, 51], and even how it was cleverly implemented to take advantage of the fact it is operating on

a TFR, it never the less has issues with constant tones, i.e., perfectly horizontal lines in the TFR that result from a constant frequency existing for some non-negligible time duration. This issue was identified after seeing repeated detections of multiple ridges when one ridge was the expected results. The issue was then confirmed by calculating the power spectral density (PSD) of the signal whose WVD result was resulting in the detection of two ridges. Because the input signal was a constant frequency sinusoid the presence of only one frequency is easily identified in both the PSD and in the WVD's results. By examining the detected ridges it was also clear that the “two” ridges are segments of the same ridge but with a non-negligible gap between the two segments.

Ultimately, either a new ridge detector algorithm or a careful inspection and testing of the CC software implementation [81] is needed. One of the experimental ridge detectors provided by Carmona, that utilizes a principle component analysis approach to detecting energy concentration, was tested and performed well. However, due to the fact that it is not documented or published in any literature, including the very book it was supplied with, makes its pedigree less than reliable.

It also appears that the geographic information systems (GIS) domain may yield more useful ridge detection algorithms than the image processing domain. GIS use ridge detectors to analysis satellite photography for automatic road detection and extraction including intersections. This problem deals with many similar issues such as chaining together segments that should be contiguous and handling sensitivities to intensity which correspond to amplitude variation in the TFR realm [82].

Bibliography

- [1] P. Flandrin, “Some features of time-frequency representations of multicomponent signals,” in *Acoustics, Speech, and Signal Processing, IEEE International Conference on (ICASSP '84)*, vol. 9, Mar 1984, pp. 266–269.
- [2] L. Cohen, “Time-frequency distributions - a review,” *Proceedings of the IEEE*, vol. 77, no. 7, pp. 941–981, Jul 1989.
- [3] L. Stanković, “A method for time-frequency analysis,” *Signal Processing, IEEE Transactions on*, vol. 42, no. 1, pp. 225–229, Jan 1994.
- [4] E. Sejdić, I. Djurović, and J. Jiang, “Time-frequency feature representation using energy concentration: An overview of recent advances,” *Digital Signal Processing*, vol. 19, no. 1, pp. 153 – 183, January 2009.
- [5] I. Daubechies, J. Lu, and H.-T. Wu, “Synchrosqueezed wavelet transforms: An empirical mode decomposition-like tool,” *Applied and Computational Harmonic Analysis*, vol. 30, no. 2, pp. 243 – 261, March 2011.
- [6] I. Daubechies, *Ten Lectures on Wavelets*, ser. CBMS-NSF Regional Conference Series in Applied Mathematics. Philadelphia, PA, USA: Society for Industrial Applied Mathematics (SIAM), 1992, ISBN: 978-0-89871-274-2.
- [7] T. A. Claasen and W. Mecklenbräuker, “The wigner distribution - a tool for time-frequency signal analysis - part 1 - continuous-time signals,” *Philips Journal of Research*, vol. 35, no. 3, pp. 217–250, 1980.
- [8] R. L. Allen and D. W. Mills, *Signal Analysis: Time, Frequency, Scale, and Structure*. Hoboken, NJ, USA: Wiley-IEEE, June 2004.

- [9] A. J. E. M. Janssen, “Positivity and spread of bilinear time-frequency distributions,” in *The Wigner Distribution: Theory and Applications in Signal Processing*, W. Mecklenbräuker and F. Hlawatsch, Eds. Amsterdam: Elsevier, 1997, ch. 1, pp. 1–58, ISBN: 0-444-88856-X.
- [10] E. Wigner, “On the quantum correction for thermodynamic equilibrium,” *Phys. Rev.*, vol. 40, pp. 749–759, Jun 1932.
- [11] J.-A. Ville, “Théorie et applications de la notion de signal analytique,” *Câbles et Transmissions*, vol. 1, pp. 61–74, 1948.
- [12] E. Boashash and et al., *Time-Frequency Signal Analysis and Processing-A Comprehensive Reference*, B. Boashash, Ed. Oxford: Elsevier Science, October 2003.
- [13] B. Boashash, “Time frequency signal analysis: Past, present and future trends,” in *Digital Control and Signal Processing Systems and Techniques*, ser. Control and Dynamic Systems, C. T. Leondes, Ed. Academic Press, 1996, vol. 78, pp. 1 – 69.
- [14] —, “Note on the use of the wigner distribution for time-frequency signal analysis,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 36, no. 9, pp. 1518–1521, Sep 1988.
- [15] L. Cohen, “Generalized phase-space distribution functions,” *Journal of Mathematical Physics*, vol. 7, no. 5, pp. 781–786, May 1966.
- [16] P. Flandrin, “Time frequency and chirps,” in *Wavelet Applications VIII*, H. H. S. an David L. Donoho, A. W. Lohmann, W. J. Campbell, and J. R. Buss, Eds., vol. 4391, Orlando, FL, Apr 2001, pp. 161–175.
- [17] L. Debnath, “Recent developments in the wigner-ville distribution and time-frequency signal analysis,” *Proceedings of the Indian National Science Academy, Part A*, vol. 68, no. 1, pp. 35–56, Jan 2002.

- [18] J. O'Toole, M. Mesbah, and B. Boashash, "A new discrete analytic signal for reducing aliasing in the discrete wigner-ville distribution," *Signal Processing, IEEE Transactions on*, vol. 56, no. 11, pp. 5427–5434, Nov 2008.
- [19] T. A. Claasen and W. Mecklenbräuker, "The wigner distribution - a tool for time-frequency signal analysis - part 2 - discrete-time signals," *Philips Journal of Research*, vol. 35, no. 4-5, pp. 276–300, 1980.
- [20] —, "The wigner distribution - a tool for time-frequency signal analysis - part 3 - relations with other time-frequency signal transformations," *Philips Journal of Research*, vol. 35, no. 6, pp. 372–389, 1980.
- [21] P. Flandrin, *Time-Frequency/Time-Scale Analysis*, ser. Wavelet Analysis and Its Applications. Academic Press, 1999, vol. 10.
- [22] L. Cohen, *Time-Frequency Analysis*, ser. Prentice Hall Signal Processing Series. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [23] F. Hlawatsch and G. Boudreaux-Bartels, "Linear and quadratic time-frequency signal representations," *Signal Processing Magazine, IEEE*, vol. 9, no. 2, pp. 21–67, April 1992.
- [24] F. Hlawatsch, "Interference terms in the wigner distribution," in *International Conference on Digital Signal Processing*, vol. 84, Florence, Italy, Sept 1984, pp. 363–367.
- [25] P. Flandrin and B. Escudié, "An interpretation of the pseudo-wigner-ville distribution," *Signal Processing*, vol. 6, no. 1, pp. 27 – 36, Jan 1984.
- [26] B. Boashash, "Theory, implementation and application of time-frequency signal analysis using the wigner-ville distribution," *Journal of Electrical and Electronics Engineering, Australia*, vol. 7, no. 3, pp. 166–177, Jul 1987.

- [27] H.-I. Choi and W. J. Williams, “Improved time-frequency representation of multicomponent signals using exponential kernels,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 37, no. 6, pp. 862–871, Jun 1989.
- [28] L. Cohen, “What is a multicomponent signal?” in *Acoustics, Speech, and Signal Processing, 1992. ICASSP-92., 1992 IEEE International Conference on*, vol. 5, Mar 1992, pp. 113–116.
- [29] L. Cohen and C. Lee, “Instantaneous frequency, its standard deviation and multicomponent signals,” in *Advanced Algorithms and Architectures for Signal Processing III*, F. T. Luk, Ed., vol. 0975, San Diego, CA, USA, August 1988, pp. 186–208.
- [30] F. Hlawatsch and P. Flandrin, “The interference structure of the wigner distribution and related time-frequency signal representations,” in *The Wigner Distribution: Theory and Applications in Signal Processing*, W. Mecklenbräuker and F. Hlawatsch, Eds. Amsterdam: Elsevier, 1997, ch. 2, pp. 59–133, ISBN: 0-444-88856-X.
- [31] J. M. O’Toole, “Discrete quadratic time-frequency distributions: Definition, computation, and a newborn electroencephalogram application,” Doctoral Dissertation, University of Queensland, Mar 2009, major Advisor: Boualem Boashash.
- [32] J. M. O’Toole, M. Mesbah, and B. Boashash, “Proofs for discrete time-frequency distribution properties,” University of Queensland, Perinatal Research Centre, Technical Report UQCCR-2009-06-09, June 2009.
- [33] J. O’Toole, M. Mesbah, and B. Boashash, “A computationally efficient implementation of quadratic time-frequency distributions,” in *9th International Symposium on Signal Processing and Its Applications (ISSPA 2007)*, vol. 1, Sharjah, United Arab Emirates, Feb 12-15 2007, pp. 288–291.

- [34] J. O’Toole and B. Boashash, “Fast and memory-efficient algorithms for computing quadratic time-frequency distributions,” *Applied and Computational Harmonic Analysis*, vol. 35, no. 2, pp. 350 – 358, Sept 2013.
- [35] J. M. O’Toole. (2015, Nov) Algorithms for (quadratic) timefrequency distributions. Website. http://otoolej.github.io/code/fast_TFDs/.
- [36] J. M. O’Toole and B. Boashash, “Memory-efficient algorithms for quadratic TFDs,” in *Time-Frequency Signal Analysis and Processing: A Comprehensive Reference*, 2nd ed., B. Boashash, Ed. Academic Press, Dec 2015, ch. 6.6, pp. 374–385.
- [37] J. M. O’Toole. (2015, Nov) Fast and memory efficient algorithms for timefrequency distributions. Website. http://otoolej.github.io/code/memeff_TFDs/.
- [38] —, Personnal Communcation, Aug 2016.
- [39] F. Auger and Éric Chassande-Mottin, “Quadratic time-frequency analysis i: Cohens class,” in *Time-Frequency Analysis: Concepts and Methods*, F. Hlawatsch and F. Auger, Eds. London, UK: ISTE, Aug 2008, ch. 5.
- [40] L. Stanković, M. Daković, and T. Thayaparan, *Time-Frequency Signal Analysis with Applications*. Norwood, MA, USA: Artech House, 2013, ISBN: 1608076520.
- [41] A. Omidvarnia, G. Azemi, J. M. O’Toole, and B. Boashash, “Robust estimation of highly-varying nonlinear instantaneous frequency of monocomponent signals using a lower-order complex-time distribution,” *Signal Processing*, vol. 93, no. 11, pp. 3251 – 3260, 2013.
- [42] S. Kadambe and G. Boudreaux-Bartels, “A comparison of the existence of “cross terms” in the wigner distribution and the squared magnitude of the wavelet transform and the short-time fourier transform,” *Signal Processing, IEEE Transactions on*, vol. 40, no. 10, pp. 2498–2517, Oct 1992.

- [43] I. Djurović and L. Stanković, “A virtual instrument for time-frequency analysis,” *Instrumentation and Measurement, IEEE Transactions on*, vol. 48, no. 6, pp. 1086–1092, Dec 1999.
- [44] M. V. Berry, “Semi-classical mechanics in phase space: A study of wigner’s function,” *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 287, no. 1343, pp. 237–271, Oct 1977.
- [45] P. Flandrin and F. Hlawatsch, “Signal representation geometry and catastrophes in the time-frequency plane,” in *Mathematics in Signal Processing*, T. Durrani, J. Abbiss, J. Hudson, R. Madan, J. McWorter, and T. Moore, Eds. Oxford, UK: Clarendon Press, 1987, pp. 3–14, (Proc. IMA Conf. Mathematics in Signal Processing, Bath, UK, Sept. 1985).
- [46] S. Stanković, I. Orović, and E. Sejdić, *Multimedia Signals and Systems*, 1st ed. Springer US, 2012, this is NOT L. Stankovic.
- [47] S. Mallat, *Wavelet Tour of Signal Processing*, 2nd ed. San Diego, CA: Academic Press, Sept 1999, ch. 4, pp. 67–124, ISBN: 9780124666061.
- [48] G. R. Putland and B. Boashash, “Can a signal be both monocomponent and multicomponent?” in *Third Australasian Workshop on Signal Processing Applications (WoSPA 2000)*, Brisbane, Australia, Dec 2000.
- [49] F. Auger, P. Flandrin, Y.-T. Lin, S. McLaughlin, S. Meignen, T. Oberlin, and H.-T. Wu, “Time-frequency reassignment and synchrosqueezing: An overview,” *Signal Processing Magazine, IEEE*, vol. 30, no. 6, pp. 32–41, Nov 2013.
- [50] B. Boashash, N. A. Khan, and T. Ben-Jabeur, “Time-frequency features for pattern recognition using high-resolution TFDs: A tutorial review,” *Digital Signal Processing*, vol. 40, pp. 1 – 30, May 2015.

- [51] R. Carmona, W. Hwang, and B. Torresani, “Multiridge detection and time-frequency reconstruction,” *Signal Processing, IEEE Transactions on*, vol. 47, no. 2, pp. 480–492, Feb 1999.
- [52] R. Carmona, W.-L. Hwang, and B. Torr  sani, *Practical Time-Frequency Analysis: Gabor and Wavelet Transforms with an Implementation in S*. Academic Press, 1998, ISBN: 978-0-12-160170-6.
- [53] B. Ristic and B. Boashash, “Scale domain analysis of a bat sonar signal,” in *Time-Frequency and Time-Scale Analysis, 1994., Proceedings of the IEEE-SP International Symposium on*, Oct 1994, pp. 373–376.
- [54] E. Boashash and et al., *Time-Frequency Signal Analysis and Processing: A Comprehensive Reference*, 2nd ed., B. Boashash, Ed. Academic Press, 2015.
- [55] J. O’Toole, M. Mesbah, and B. Boashash, “Improved discrete definition of quadratic time-frequency distributions,” *Signal Processing, IEEE Transactions on*, vol. 58, no. 2, pp. 906–911, Feb 2010.
- [56] S. G  mez, V. Naranjo, and R. Miralles, “Removing interference components in time-frequency representations using morphological operators,” *Journal of Visual Communication and Image Representation*, vol. 22, no. 5, pp. 401–410, July 2011.
- [57] M. Davies and C. James, “Source separation using single channel ica,” *Signal Processing*, vol. 87, no. 8, pp. 1819 – 1832, August 2007.
- [58] Z. Yu, Y. Sun, and W. Jin, “A novel generalized demodulation approach for multi-component signals,” *Signal Processing*, vol. 118, pp. 188 – 202, Jan 2016.
- [59] C. Jutten and J. Herault, “Blind separation of sources, part I: An adaptive algorithm based on neuromimetic architecture,” *Signal Processing*, vol. 24, no. 1, pp. 1 – 10, July 1991.

- [60] D. Barry, E. Coyle, D. Fitzgerald, and R. Lawlor, “Single channel source separation using short-time independent component analysis,” in *Audio Engineering Society Convention 119*, New York, New York, USA, Oct 2005.
- [61] X. Chen, A. Liu, H. Peng, and R. K. Ward, “A preliminary study of muscular artifact cancellation in single-channel eeg,” *Sensors*, vol. 14, no. 10, pp. 18 370–18 389, 2014.
- [62] C. James and D. Lowe, “Single channel analysis of electromagnetic brain signals through ica in a dynamic systems framework,” in *Engineering in Medicine and Biology Society, 2001. Proceedings of the 23rd Annual International Conference of the IEEE*, vol. 2, Istanbul, Turkey, Oct 2001, pp. 1974–1977 vol.2.
- [63] N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, N.-C. Yen, C. C. Tung, and H. H. Liu, “The empirical mode decomposition and the hilbert spectrum for nonlinear and non-stationary time series analysis,” *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 454, no. 1971, pp. 903–995, Mar 1998.
- [64] G. Rilling, P. Flandrin, and P. Gonçalvès, “On empirical mode decomposition and its algorithms,” in *Proceedings of the 6th IEEE/EURASIP Workshop on Nonlinear Signal and Image Processing (NSIP '03)*, Grado, Italy, June 2003.
- [65] G. Rilling and P. Flandrin, “One or two frequencies? the empirical mode decomposition answers,” *Signal Processing, IEEE Transactions on*, vol. 56, no. 1, pp. 85–95, Jan 2008.
- [66] Z. Wu and N. E. Huang, “Ensemble empirical mode decomposition: A noise-assisted data analysis method,” *Advances in Adaptive Data Analysis*, vol. 1, no. 1, pp. 1–41, Jan 2009.

- [67] X. Hu, S. Peng, and W. L. Hwang, “Emd revisited: A new understanding of the envelope and resolving the mode-mixing problem in am-fm signals,” *IEEE Transactions on Signal Processing*, vol. 60, no. 3, pp. 1075–1086, Mar 2012.
- [68] H.-T. Wu, P. Flandrin, and I. Daubechies, “One or two frequencies? the synchrosqueezing answers,” *Advances in Adaptive Data Analysis*, vol. 3, no. 1-2, pp. 29–39, 2011.
- [69] T. Oberlin, S. Meignen, and V. Perrier, “An alternative formulation for the empirical mode decomposition,” *Signal Processing, IEEE Transactions on*, vol. 60, no. 5, pp. 2236–2246, May 2012.
- [70] M. A. Colominas, G. Schlotthauer, and M. E. Torres, “An unconstrained optimization approach to empirical mode decomposition,” *Digital Signal Processing*, vol. 40, pp. 164–175, May 2015.
- [71] M. Torres, M. Colominas, G. Schlotthauer, and P. Flandrin, “A complete ensemble empirical mode decomposition with adaptive noise,” in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*, Prague, May 2011, pp. 4144–4147.
- [72] M. A. Colominas, G. Schlotthauer, and M. E. Torres, “Improved complete ensemble emd: A suitable tool for biomedical signal processing,” *Biomedical Signal Processing and Control*, vol. 14, pp. 19 – 29, Nov 2014.
- [73] S. Meignen and V. Perrier, “A new formulation for empirical mode decomposition based on constrained optimization,” *IEEE Signal Processing Letters*, vol. 14, no. 12, pp. 932–935, Dec 2007.
- [74] S. H. Mankad and S. N. Pradhan, “Application of software defined radio for noise reduction using empirical mode decomposition,” in *2nd International Conference on Computer Science, Engineering and Applications (ICCSEA 2012)*, C. D.

- Wyld, J. Zizka, and D. Nagamalai, Eds. New Delhi, India: Springer Berlin Heidelberg, May 25-27 2012, pp. 113–121.
- [75] T. Kijewski-Correa and A. Kareem, “Nonlinear signal analysis: Time-frequency perspectives,” *Journal of Engineering Mechanics*, vol. 133, no. 2, pp. 238–245, 2007.
- [76] D. P. Mandic, N. u. Rehman, Z. Wu, and N. E. Huang, “Empirical mode decomposition-based time-frequency analysis of multivariate signals: The power of adaptive data analysis,” *Signal Processing Magazine, IEEE*, vol. 30, no. 6, pp. 74–86, Nov 2013.
- [77] N. E. Huang, M.-L. C. Wu, S. R. Long, S. S. Shen, W. Qu, P. Gloersen, and K. L. Fan, “A confidence limit for the empirical mode decomposition and hilbert spectral analysis,” *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 459, no. 2037, pp. 2317–2345, Sept 2003.
- [78] N. E. Huang and Z. Wu, “A review on hilbert-huang transform: Method and its applications to geophysical studies,” *Reviews of Geophysics*, vol. 46, no. 2, pp. 23–pgs., 2008.
- [79] A. Papandreou-Suppappola, F. Hlawatsch, and G. Boudreaux-Bartels, “Quadratic time-frequency representations with scale covariance and generalized time-shift covariance: A unified framework for the affine, hyperbolic, and power classes,” *Digital Signal Processing*, vol. 8, no. 1, pp. 3 – 48, Jan 1998.
- [80] M. A. Awal and B. Boashash, “An automatic fast optimization of quadratic time-frequency distribution using the hybrid genetic algorithm,” *Signal Processing*, vol. 131, pp. 134 – 142, Feb 2017.
- [81] R. Carmona. S-wave software. Available: <https://www.princeton.edu/~rcarmona/mainpublications.html>.

- [82] C. Steger, “An unbiased detector of curvilinear structures,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 20, no. 2, pp. 113–125, Feb 1998.

Appendix

Appendix A

Software Resources

The following software resources were utilized in the creation of the software implementation of the mono-component measure.

- [1] Mathworks. Matlab. Commercial Software. [Online]. Available: www.mathworks.com
- [2] J. M. O’Toole. Fast and Memory-Efficient Algorithms for Quadratic Time-Frequency Distributions. [Online]. Available: https://github.com/otoolej/memeff_TFDs
- [3] R. Carmona. S-wave. [Online]. Available: <https://www.princeton.edu/~rcarmona/mainpublications.html>
- [4] G. Rilling and P. Flandrin. Matlab/C codes for EMD and EEMD with examples. [Online]. Available: <http://perso.ens-lyon.fr/patrick.flandrin/emd.html>
- [5] F. Auger, P. Flandrin, P. Gonçalves, and O. Lemoine. Time-Frequency Toolbox. [Online]. Available: <http://tftb.nongnu.org/>

Vita

Austin P. Albright is the son of Steve and Peggy Albright. In May of 2004, he graduated from Tennessee Technological University (TTU) in Cookeville, Tennessee with a Bachelor of Science in Electrical Engineering. Upon graduating, he began working full-time at Oak Ridge National Laboratory (ORNL) through a student research experience program till entering graduate school full-time in 2006.

In 2005, Austin began working on a Master of Science in Electrical Engineering at the University of Tennessee - Knoxville (UTK). During his graduate studies Austin was awarded a prestigious Department of Homeland Security Fellowship in a national competition to become the only Department of Homeland Security Fellow in the entire state of Tennessee. In 2007, Austin was inducted into the Beta Phi chapter of Eta Kappa Nu, the Electrical Engineering honor society, at the University of Tennessee.

He graduated with a Master's degree in Electrical Engineering in December of 2007, after defending his Master's thesis titled, "*The Detection of Stress Corrosion Cracking in Natural Gas Pipelines Using Electromagnetic Acoustic Transducers*" under the supervision of Dr. Hairong Qi. Upon graduating with his Master's degree he was hired full-time as a research and development engineer at ORNL in the Cognitive Radio group of the Radio Frequency, Communications, and Intelligent Systems (RFCIS) Group. Where he has specialized in Software Defined Radio (SDR), specifically the implementation of SDRs in Field Programmable Gate Arrays (FPGAs), hardware-software interfaces, and signal analysis and processing.

In 2014, thanks to receiving an educational sabbatical from ORNL, Austin was

able to return to UTK to complete the remaining requirements to earn a Ph.D. in Electrical Engineering under the supervision of Dr. Hairong Qi. In 2015, Austin was honored to be elected to the grade of Senior Member of the Institute of Electrical and Electronics Engineers (IEEE). He continues to work in the fields of SDR and signal processing at ORNL in the RFCIS group. Austin currently lives in Knoxville, TN with his wife Melissa and their four children.