Wavelets

Algorithms & Applications

Wavelets

Algorithms & Applications

Yves Meyer
CEREMADE
and
Institut Universitaire de France

Translated and Revised by

Robert D. Ryan
Office of Naval Research
European Office

Society for Industrial and Applied Mathematics
Philadelphia 1993

siam.

Library of Congress Cataloging-in-Publication Data

Meyer, Yves.
[Ondelettes. English]
Wavelets : algorithms and applications / Yves Meyer ; translated by Robert D. Ryan.
p. cm.
"Translation based on lectures originally presented by the author for the Spanish Institute in Madrid, Spain, February 1991"—CIP galley.
ISBN 0-89871-309-9
1. Wavelets. I. Title.
QA403.3.M4913 1993
515'.2433—dc20 93-15100

Translator's Foreword

The question most often asked by those who heard I was translating this book was, "How did you get the job?" Well, I asked for it. I mentioned to Professor Meyer in March 1992 that I had heard that he had written a new book. He said, yes, and that the book was based on notes from lectures he had given at the Spanish Institute in Madrid. He added that the book was being translated into Spanish and that SIAM was interested in publishing an English edition, for which they would need a translator. I volunteered to do the job; what you have is the result.

In addition to translating the text, I have tried to "work through" much of the mathematics to correct typos. I have also added a line of explanation here and there where it seemed appropriate, and sections of text and references have been updated. These revisions are not highlighted in any particular way (i.e., there are no translator's notes except at one reference), but rather incorporated into the text. These changes were made with Professor Meyer's blessing. Of course, there is always the possibility that in the process of updating the manuscript I have introduced other errors; for these I take full responsibility.

The great fun of this project has been the chance to work with Yves Meyer and other members of "team wavelet." Professor Meyer improved both my French and my mathematics, and his enthusiasm and appreciation for my efforts kept things moving. Direct help also came from John Benedetto, Marie Farge, Patrick Flandrin, Stéphane Jaffard, and Hamid Krim. Alex Grossmann gave moral support by assuring me that it was an important project. My sincere thanks to all of these people. The work was done while I was a Liaison Scientist for the Office of Naval Research European Office, where my primary job was to report on mathematics in Europe. It was through this work that I first made contact with the French wavelet community, and I thank the Office of Naval Research for that opportunity. This was essentially a weekend and evening project, and hence a family project. In this context, I thank my son, Michael J. Ryan, who kept house and produced great dinners while I kept the electrons moving.

Robert D. Ryan
November 1992
London, UK

Preface

The "theory of wavelets" stands at the intersection of the frontiers of mathematics, scientific computing, and signal processing. Its goal is to provide a coherent set of concepts, methods, and algorithms that are adapted to a variety of nonstationary signals and that are also suitable for numerical signal processing.

This book results from a series of lectures that Mr. Miguel Artola Gallego, Director of the Spanish Institute, invited me to give on wavelets and their applications. I have tried to fulfill, in the following pages, the objective the Spanish Institute set for me: to present to a scientific audience coming from different disciplines, the prospects that wavelets offer for signal and image processing.

A description of the different algorithms used today under the name "wavelets" (Chapters 2–7) will be followed by an analysis of several applications of these methods: to numerical image processing (Chapter 8), to fractals (Chapter 9), to turbulence (Chapter 10), and to astronomy (Chapter 11). This will take me out of my scientific domain; as a result, the last two chapters are merely resumes of the original articles on which they are based.

I wish to thank the Spanish Institute for its generous hospitality as well as its Director for his warm welcome. Additionally, I note the excellent organization by Mr. Pedro Corpas.

My thanks go also to my Spanish friends and colleagues who took the time to attend these lectures.

Contents

CHAPTER 1

Signals and Wavelets

The purpose of this first chapter is to give the reader a fairly clear idea about the scientific content of the following chapters. All of the "themes" that will be developed in this study, using the inevitable mathematical formalism, already appear in this "overture." It is written with a concern for simplicity and clarity, while avoiding as much as possible the use of formulas and symbols.

Signal and image processing always leads to a collection of techniques or procedures. But like all other scientific disciplines, signal and image processing assumes certain preliminary scientific conventions. We have sought, in this first chapter, to describe the intellectual architecture underlying the algorithmic constructions that will be presented in the following chapters.

1.1. What is a signal?

Signal processing has become an essential part of contemporary scientific and technological activity. Signal processing is used in telecommunications (telephone and television), in the transmission and analysis of satellite images, and in medical imaging (echography, tomography, and nuclear magnetic resonance), all of which involve the analysis and interpretation of complex time series. A record of stock price fluctuations is a signal, as is a record of temperature readings that permit the analysis of climatic variations and the study of global warming. This list is by no means exhaustive.

Does there exist a precise definition of a signal that is appropriate for the field of scientific activity called "signal processing"? A needlessly broad definition could include the sequence of letters, spaces, and punctuation marks appearing in Montaigne's *Essays*, but the tools we present do not apply to such a signal. However, the structuralist analysis done by Roland Barthes on literary texts shares some amusing similarities with the multiresolution analysis that we describe in Chapter 4.

The signals we study will always be series of numbers and not series of letters, words, or phrases. These numbers come from measurements, which are typically made using some recording method. The signals ultimately appear as functions of time. This is true for one-dimensional signals. The case of two-dimensional signals will be examined in a moment.

The objectives of signal processing are to analyze accurately, code efficiently, transmit rapidly, and then to reconstruct carefully at the receiver the delicate *oscillations or fluctuations of this function of time. This is important because all of the information contained in the signal is effectively present and hidden in the complicated arabesques appearing in its graphical representation.*

These remarks apply to speech: A speech signal originates as subtle time variations of air pressure and becomes a curve whose complex graphical characteristics are an "adapted copy" of the voice.

It is equally important to consider two-dimensional signals, which is to say, images. Here again, image processing is done on the numerical representation of the image. For a black and white image, the numerical representation is created by replacing the x and y coordinates of an image point with those of the closest point on a sufficiently fine grid. The value $f(x, y)$ of the "gray scale" is then replaced with an average coefficient, which is then assigned to the corresponding grid point.

The image thus becomes a large, typically square, matrix. Image processing is done on this matrix.

These matrices are enormous, and as soon as one deals with a sequence of images, the volume of numerical data that must be processed becomes immense. Is it possible to reduce this volume by considering the "hidden laws" or correlations that exist among the different pieces of numerical information representing the image? This question leads us naturally to define the goals of the scientific discipline called "signal processing."

1.2. The goals of signal and image processing.

Experts in signal processing are called on to describe, for a given class of signals, algorithms that lead to the construction of microprocessors and that allow certain operations and tasks to be done automatically. These tasks may be: *analysis and diagnostics*, *coding*, *quantization and compression*, *transmission or storage*, *and synthesis and reconstruction*.

We will use several examples to illustrate the nature of these operations and the difficulties they present. It will become clear that no "universal algorithm" is appropriate for the extreme diversity of the situations encountered. Thus, a large part of this work is devoted to constructing coding or analysis algorithms that can be adapted to the signals that one processes.

Our first example is the study of climatic variations and global warming. This example was discussed by Professor Jacques-Louis Lions at the Spanish Institute in 1990, and the following thoughts were inspired by his talks [4].

In this example, one has fairly precise temperature measurements from different points in the northern hemisphere that were taken over the last two centuries, and one tries to discover if industrial activity has caused global warming. The extreme difficulty of the problem arises from the existence of significant natural temperature fluctuations. Moreover, these fluctuations and the corresponding climatic changes have always existed, as we learn from paleoclimatology [6].

To specify a *diagnostic*, it is essential to *analyze*, and then to *erase*, these natural fluctuations (which play the role of noise) in order to have access to the "artificial" heating of the planet resulting from human activity. The diagnostic often depends on extracting a small number of significant parameters from a signal whose complexity and size are overwhelming. Thus the analysis and the diagnostic rely naturally on *data compression.* If this compression is done inappropriately, it can falsify the diagnostic.

Data compression also occurs in the problem of *transmission.* Indeed, transmission channels have a limited capacity, and it is therefore important to reduce as much as possible the abundance of raw information so that it fits within the channel's "bit allocation."

One thinks, for example, of the digital telephone (Chapter 3) and the 64 Kbit/sec standard, which limits, without appeal, the quantity of information that can be transmitted in one second.

A more surprising example appears in neurophysiology. The optic nerve's capacity to transmit visual information is clearly *less* than the volume of information collected by all the retinal cells. Thus, there must be "low-level processing" of information before it transits the optic nerve. David Marr has developed a theory that allows us to understand the purpose and performance of this low-level processing. We present this theory in Chapter 8.

We now consider problems posed by *coding* and *quantization.* Different coding algorithms will be presented and studied in this work: subband coding, transform coding, and coding by zero-crossings. In each case, coding involves methods to transform the recorded numerical signal into another representation that is, depending on the nature of the signals studied, more convenient for some task or further processing. Quantization is associated with coding. The "exact" numerical values given by coding are replaced with nearby values that are compatible with the bit allocation dictated by the transmission capacity.

Quantization is an unavoidable step in signal and image processing. Unfortunately, it introduces systematic errors, known as "quantization noise." The coding algorithms that are used (taking into account the nature of the signals) ought to reduce the effects of quantization noise when decoding takes place. One of the advantages of *quadrature mirror filters* is that they "trap" this quantization noise inside well-defined frequency channels. These filters will be studied in Chapter 3.

The problems encountered in *archiving* data (as well as problems of transmission and reconstruction) are illustrated by the FBI's task of storing the American population's fingerprints. Different image-compression algorithms were tested, and a variant of the algorithm described in Chapter 6 gave the best results. This established a standard for fingerprint compression and reconstruction.

The last group of operations consists of *decoding, synthesis,* and *restoration.* Synthesis and decoding are the inverse operations of coding and quantization. The task is to reconstruct an image or audible signal at the receiver from the series of 0's and 1's that have traveled over the transmission channel. One thinks of decoding an encoded message, as in cryptography, and this analogy is correct because one cannot reconstruct an image or signal without knowing the coding algorithm.

Signal restoration is similar to the restoration of old paintings. It amounts to ridding the signal of artifacts and errors (which we call noise), and to enhancing certain aspects of the signal that have undergone attenuation, deterioration, or degradation.

1.3. Stationary signals, transient signals, and adaptive coding algorithms.

We have just defined a set of tasks, or operations, to be performed on signals or images. These tasks form a coherent collection. The purpose of this book is to describe certain coding algorithms that have, during the last few years, been shown to be particularly effective for analyzing signals having a fractal structure or for compression and storage. We will also describe certain "meta-algorithms" that allow one to choose the coding algorithm best suited to a given signal. To better approach this problem of choosing an adaptive algorithm, we briefly classify signals by distinguishing stationary signals, quasi-stationary signals, and transient signals.

A signal is stationary if its properties are statistically invariant over time. A well-known stationary signal is white noise, which, in its sampled form, appears as a series of independent drawings. A stationary signal can exhibit unexpected events, but we know in advance the probabilities of these events. These are the statistically predictable unknowns.

The ideal tool for studying stationary signals is the Fourier transform. In other words, stationary signals decompose canonically into linear combinations of waves (sines and cosines). In the same way, signals that are not stationary decompose into linear combinations of wavelets.

The study of nonstationary signals, where transient events appear that cannot be predicted (even statistically with knowledge of the past), necessitates techniques different from Fourier analysis. These techniques, which are specific to the nonstationarity of the signal, include wavelets of the "time-frequency" type and wavelets of the "time-scale" type. "Time-frequency" wavelets are suited, most specifically, to the analysis of quasi-stationary signals, while "time-scale" wavelets are adapted to signals having a fractal structure.

Before defining "time-frequency" wavelets and "time-scale" wavelets, we indicate their common points. They belong to a more general class of algorithms that are encountered as often in mathematics as in speech processing. Mathematicians speak of "atomic decompositions," while speech specialists speak of "decompositions in time-frequency atoms"; *the scientific reality is the same in both cases.*

An "atomic decomposition" consists in extracting the simple constituents that make up a complicated mixture. However, contrary to what happens in chemistry, the "atoms" that are discovered in a signal will depend on the point of view adopted for the analysis. These "atoms" will be "time-frequency atoms" when we study quasi-stationary signals, but they could, in other situations, be replaced by "time-scale wavelets" or "Grossmann–Morlet wavelets."

These "atoms" or "wavelets" have no more physical existence than the number system used to multiply the mass of the earth by that of the moon. Each number system has an internal logical coherence, but no scientific law asserts that multiplication must of necessity be done in base 10 rather than base 2. On the other hand, the number system used by the Romans is certainly excluded because it is not suitable for multiplication.

Having different algorithms that allow us to code a signal by decomposing it into "time-frequency atoms" presents us with a similar situation. The decision to use one or the other of these algorithms will be made by considering their *performance.* How well they perform must be judged in terms of one of the anticipated *goals* of signal processing. *An algorithm that is optimal for compression can be disastrous for analysis: A standard energetic criterion for the compression could cause details that are important for the analysis to be systematically neglected.*

These thoughts will be developed and clarified in §§1.6 and 1.7. It is time to define wavelets, which we do in the next two sections.

1.4. Grossmann–Morlet time-scale wavelets.

"Time-scale" analysis (which should be called "space-scale" in the image case, but which we prefer to call "multiresolution analysis") involves using a vast range of scales for signal analysis. This notion of scale, which clearly refers to cartography, implies that the signal (or image) is replaced, at a given scale, by the best possible approximation that can be drawn at that scale. By "traveling" from the large scales toward the fine scales, one "zooms in" and arrives at more and more exact representations of the given signal.

The analysis is then done by calculating the change from one scale to the next. These are the details that allow one, by correcting a rather crude approximation, to move toward a better quality representation. This algorithmic scheme is called "multiresolution analysis" and is developed in Chapters 3 and 4. *Multiresolution analysis is equivalent to an atomic decomposition where the atoms are Grossmann–Morlet wavelets.*

We define these wavelets by starting with a function $\psi(t)$ of the real variable t. This function is called a "mother wavelet" provided it is well localized and oscillating. (By oscillating it resembles a wave, but by being localized it is a wavelet.) The localization condition is expressed in the usual way as decreasing rapidly to zero when $|t|$ tends to infinity. The second condition suggests that $\psi(t)$ vibrates like a wave: Here we require that the integral of $\psi(t)$ be zero and that the same hold true for the first m movements of ψ. This is expressed as

$$0 = \int_{-\infty}^{\infty} \psi(t)\,dt = \cdots = \int_{-\infty}^{\infty} t^{m-1}\psi(t)\,dt. \tag{1.1}$$

The "mother wavelet," $\psi(t)$, generates the other wavelets, $\psi_{(a,b)}(t)$, $a > 0$, $b \in \mathbb{R}$, of the family by change of scale (the scale of $\psi(t)$ is conventionally 1, and that of $\psi_{(a,b)}(t)$ is $a > 0$) and translation in time (the function $\psi(t)$ is conventionally centered around 0, and $\psi_{(a,b)}(t)$ is then centered around b).

Thus we have

$$\psi_{(a,b)}(t) = \frac{1}{\sqrt{a}}\psi\left(\frac{t-b}{a}\right), \qquad a > 0, \qquad b \in \mathbb{R}. \tag{1.2}$$

Alex Grossmann and Jean Morlet have shown that, if $\psi(t)$ is real-valued, this collection can be used as if it were an orthonormal basis. This means that any signal of finite energy can be represented as a linear combination of wavelets $\psi_{(a,b)}(t)$ and that the coefficients of this combination are, up to a normalizing factor, the scalar products $\int_{-\infty}^{\infty} f(t)\psi_{(a,b)}(t)\,dt$.

These scalar products measure, in a certain sense, the fluctuations of the signal $f(t)$ around the point b, at the scale given by $a > 0$.

It required uncommon scientific intuition to assert, as Grossmann and Morlet did, that this new method of time-scale analysis was suitable for the analysis and synthesis of *transient signals*. Signal processing experts were annoyed by the intrusion of these two poachers on their preserve and made fun of their claims.

This polemic died out after only a few years. In fact, the argument should never have arisen because the methods of time-scale or multiresolution analysis had existed for five or six years under various disguises: in signal analysis (under the name of quadrature mirror filters) and in image analysis (under the name of pyramid algorithms).

The first to report on this was Stephane Mallat. He constructed a guide that allowed the same signal analysis method to be recognized under very different presentations: wavelets, pyramid algorithms, quadrature mirror filters, Littlewood–Paley analysis, David Marr's zero crossings...

Mallat's brilliant synthesis has been the source of many new developments. One of the most notable of these is Ingrid Daubechies's discovery of orthonormal wavelet bases having preselected regularity and compact support. The only previously known case was the Haar system (1909), which is not regular. Thus 80 years separated Alfred Haar's work and its natural extension by Daubechies [1].

1.5. Time-frequency wavelets from Gabor to Malvar.

Dennis Gabor was the first to introduce *time-frequency wavelets* or *Gabor wavelets*. He had the idea to divide a wave (whose mathematical representation is $\cos(\omega t + \varphi)$) into segments and then to throw away all but one of these segments. This left a piece of a wave, or a wavelet, which had a beginning and an end.

To use a musical analogy, a wave corresponds to a note (a *ré* minor, for example) that has been emitted since the origin of time and sounds indefinitely, without attenuation, until the end of time. A wavelet then corresponds to the same *ré* minor that is struck at a certain moment, say, on a piano, and is later muffled by the pedal. In other words, a Gabor wavelet has (at least) three pieces of information: a beginning, an end, and a specific frequency in between.

Difficulties appeared when it was necessary to decompose a signal using Gabor wavelets. As long as one does only continuous decompositions (using all frequencies and all time), Gabor wavelets can be used as if they formed an orthonormal basis. But the corresponding discrete algorithms do not exist, or they require so much tinkering that they become too complicated.

It is only very recently, by abandoning Gabor's approach, that two separate groups have discovered time-frequency wavelets having good algorithmic qualities. These special time-frequency wavelets, called Malvar wavelets, are particularly well suited for coding speech and music. Moreover, they are equally useful to the FBI for storing fingerprints.

The decomposition of a signal in an orthonormal basis of Malvar wavelets imitates writing music using a musical score. But this composition is misleading because a piece of music can be written in only one way, whereas there exists a nondenumerable infinity of orthonormal bases of Malvar wavelets. Choosing one among these is equivalent to segmenting the given signal and then doing a traditional Fourier analysis on the delimited pieces. What is the best way to choose this segmentation? This question leads us naturally to the following section.

1.6. Optimal algorithms in signal processing.

Which wavelet to choose? I have often posed this question at meetings on wavelets and their applications held since 1985. But this question needs to be sharpened. *What freedom of choice is at our disposal? What are the objectives of the choices we make? Can we make better use of the choices offered to us by considering the anticipated goals?* These are several of the questions we will try to answer.

The goal we have in mind is aptly illustrated by a remark by Benoît Mandelbrot made in an interview on "France-Culture": "The world around us is very complicated. The tools at our disposal to describe it are very weak."

It is notable that Mandelbrot used the word "describe" and not "explain" or "interpret." We are going to follow him in this, ostensibly, very modest approach. This is our answer to the problem about the objectives of the choices: *Wavelets, whether they are of the time-scale or time-frequency type, will not help us to explain scientific facts, but they will serve to describe the reality around us, whether or not it is scientific.*

Our task is to optimize the description. This means that we must make the best use of the resources allocated to us (for example, the number of available bits) to obtain the most precise possible description.

To resolve this problem, we must first indicate how the quality of the description will be judged. Most often, the criteria used are academic and do not correspond at all to the user's point of view.

For example, in image processing, all the calculations for judging the quality of the description use the quadratic mean value of gray levels. It is clear, however, that our eye has a much more selective sensitivity. Thus, in the last analysis, we should submit the performance of an "optimal algorithm" to the users because the average approximation criterion that leads to this algorithm will often be inadequate.

The case of speech (telephonic communication) or music is similar. After systematic research that optimizes the reception quality (quality calculated with an average), it would be advisable to experiment by taking into account the judgments of telephone users and musicians.

Thus we see a two-state program developing. Nevertheless, the only stage we describe in the following pages is the "objective search" for an optimal algorithm, even though its optimality is defined in terms of a debatable energy criterion.

Rather than formulate ad hoc algorithms for each signal or each class of signals, we are going to construct, once and for all, a vast collection called a library of algorithms. We will also construct a "meta-algorithm" whose function will be to find the particular algorithm in the library that best serves the given signal in view of the criterion for the quality of the description.

It is hardly an exaggeration to say that we will introduce almost as many analysis algorithms as there are signals. For example, for a signal recorded on $2^{10} = 1024$ points, the number of algorithms at our disposal will be of the order 2^{1024}. This is the number of all the signals defined on our 1024 points that take only two values, 0 and 1.

Thus we will use a "very large library" to describe the signals, but we exclude the "library of Babel." This would contain all the books, or all the signals in our case. And as everyone knows, the search for a specific book in the library of Babel is an insurmountable task. The "ideal library" must be sufficiently rich to suit all transient signals, but the "books" must be easily accessible.

While a single algorithm (Fourier analysis) is appropriate for all stationary signals, the transient signals are so rich and complex that a single analysis method (whether of time-scale or time-frequency) cannot serve them all.

If we stay in the relatively narrow environment of Grossmann–Morlet wavelets, also called time-scale algorithms, we have only two ways to adapt the algorithm to the signal being studied: We can choose one or another analyzing wavelet, and we can use either the continuous or the discrete version of the wavelets.

For example, we can require the analyzing wavelet ψ to be an analytic signal, which means that its Fourier transform $\hat{\psi}(\omega)$ is zero for negative frequencies. In this case, all the wavelets $\psi_{(a,b)}$, $a > 0$, $b \in \mathbb{R}$, generated by ψ will still have this property, and their linear combinations given by the algorithm will be the analytic signal F associated with the real signal f.

Likewise, we can follow Daubechies and, for a given $r \geq 1$, choose for $\psi(x)$ a real-valued function in the class C^r with compact support such that the collection $2^{j/2}\psi(2^j x - k)$, $j,k \in \mathbb{Z}$, is an orthonormal basis for $L^2(\mathbb{R})$. In this discrete version of the algorithm, $a = 2^{-j}$ and $b = k2^{-j}, j,k \in \mathbb{Z}$.

In spite of this, the choices that can be made from the set of time-scale wavelets remain limited. *The search for optimal algorithms leads us on some* *remarkable algorithmic adventures, where time-scale wavelets and time-frequency wavelets are in competition, and where they are also compared to intermediate algorithms that mix the two extreme forms of analysis.*

These considerations are developed in Chapters 6 and 7 and the question I asked myself six years ago (what wavelet to choose?) seems passé to me today. The choices that we can and must consider no longer involve only the analyzing instrument (the wavelet); they also involve the methodology employed (time-scale, time-frequency, or intermediate algorithms).

Today, the competing algorithms (time-scale and time-frequency) are included in a whole universe of intermediate algorithms. An entropy criterion permits us to choose the algorithm that optimizes the description of the given signal within the given bit-allocation.

Each algorithm is presented in terms of a particular orthogonal basis. We can compare searching for the optimal algorithm to searching for the best point of view, or best perspective, to look at a statue in a museum. Each point of view reveals certain parts of the statue and obscures others. We change our point of view to find the best one by going around the statue. In effect, we make a rotation; we change the orthonormal basis of reference to find the optimal basis.

These reflections lead us quite naturally to the scientific thoughts of David Marr, which we present in the next section.

1.7. Optimal representation, according to Marr.

David Marr was fascinated and intrigued by the complex relations that exist between the choice of a representation of a signal and the nature of the operations or transformations that such a representation permits. He wrote [5, pp. 20 21]:

A *representation* is a formal system for making explicit certain entities or types of information, together with a specification of how the system does this. And I shall call the result of using a representation to describe a given entity a description of the entity in that representation.

For example, the Arabic, Roman and binary numerical systems are all formal systems for representing numbers. The Arabic representation consists of a string of symbols drawn from the set {0, 1, 2, 3, 4, 5, 6, 7, 8, 9} and the rule for constructing the description of a particular integer n is that one decomposes n into a sum of multiples of powers of 10... A musical score provides a way of representing a symphony; the alphabet allows the construction of a written representation of words...

A representation, therefore, is not a foreign idea at all—we all use representations all the time. However, the notion that one can capture some aspect of reality by making a description of it using a symbol and that to do so can be useful seems to me a fascinating and powerful idea. But even the simple examples we have discussed introduce some rather general and important issues that arise whenever one chooses to use one particular representation.

For example, if one chooses the Arabic numerical representation, it is easy to discover whether a number is a power of 10 but difficult to discover whether it is a power of 2. If one chooses the binary representation, the situation is reversed.