Wavelets for Computer Graphics: Theory and Applications

Eric J. Stollnitz   Tony D. DeRose   David H. Salesin

Abstract

For a complete copy of the book, please visit the Morgan-Kaufmann website or Amazon.com.

Preface

Wavelets are a mathematical tool for hierarchically decomposing functions. Though rooted in approximation theory, signal processing, and physics, wavelets have also recently been applied to many problems in computer graphics. These graphics applications include image editing and compression, automatic level-of-detail control for editing and rendering curves and surfaces, surface reconstruction from contours, and fast methods for solving simulation problems in 3D modeling, global illumination, and animation.

Despite the growing evidence that wavelets are quickly becoming a core technique in computer graphics, most of the existing literature has been written primarily for the signal processing and approximation theory communities, and is relatively inaccessible to researchers working in computer graphics. In addition, most of the established theory on wavelets has been developed for the theoretically-pure case of signals of infinite length. Unfortunately, this classical theory begins to break down when it comes to representing the kinds of finite data sets—such as images, open curves, and bounded surfaces—that arise most commonly in computer graphics.

This monograph is intended to address both of these problems.

First, it should provide the computer graphics professional and researcher with a firm understanding of the theory and applications of wavelets. The reader is assumed to have had a first course in linear algebra—but to have forgotten most of it (the text includes a linear algebra refresher with most of the necessary background included. We have intentionally kept the style relatively light in tone, stressing intuition and clarity over rigor.

Second, the monograph also takes a significantly different approach from existing texts, in that it focuses on a more generalized theory of wavelets, which allows wavelets to be constructed naturally on the kinds of bounded domains that arise quite commonly in computer graphics applications. This more generalized theory, it turns out, is intimately tied to the process of recursive subdivision. Thus, the monograph also develops the theory of subdivision curves and surfaces in its treatment of wavelets.

This monograph is by no means exhaustive, neither in its development of the theory of wavelets, nor in its coverage of their applications in computer graphics. Our goal in writing this text, rather, has been to emphasize the aspects of the theory that have already proven themselves to be most useful in computer graphics, and to provide a small but broad enough set of applications to illustrate how the theory can be applied in practice—in a surprisingly wide variety of domains.

How this book came about

Our interest in wavelets began in late 1992 and early 1993. At that time wavelets had not yet (to our knowledge) been applied to problems in computer graphics, but we felt that there might be many such applications. The easiest way to learn a new area is to hold a graduate seminar on the topic, so that's what we did in Spring 1993. Within two or three weeks it became apparent to us that wavelets would have widespread use in computer graphics and geometric modeling. By the end of the quarter four students were off and running on Ph.D. topics (Per Christensen, Adam Finkelstein, Michael Lounsbery, and David Meyers), and one on a Master's Project (Debbie Berman). Much of this work was later published in the proceedings of the annual SIGGRAPH and Graphics Interface conferences, as well as in recent issues of *IEEE Computer Graphics and Applications* and *ACM Transactions on Graphics*.

This monograph is also a direct outgrowth of that seminar. Although it draws heavily on our previously published work, we have also added a considerable amount of new material, and we have attempted to develop a unified framework and consistent notation for all of the results.

Organization of the book

The book is organized into four main parts according to application area: images, curves, surfaces, and physical simulation. Within these parts, chapters on theory are intermingled with chapters on specific applications. We felt that such an organization would provide a more interesting narrative—and would also better motivate the theory—than segregating all the theory chapters from the "practice."

While we would not recommend reading the applications chapters without the chapters on theory, it should be possible to do the reverse. Thus, readers who are interested in learning about just the theory of wavelets may wish to focus their attention on Chapters 2, 3, 6, 7 and 10.

Acknowledgements

We are grateful to Sean Anderson, Jason Bartell, Debbie Berman, Per Christensen, Michael Cohen, Adam Finkelstein, Steven Gortler, Hugues Hoppe, Chuck Jacobs, Dani Lischinski, Michael Lounsbery, Dave Meyers, and Joe Warren for contributing material for the book. We are also indebted to the other members of the graphics research group at the University of Washington's Department of Computer Science and Engineering who helped make this book possible.

Thanks also to Ronen Barzel, Michael Shantzis, and Andy van Dam for reviewing preliminary chapters, and to Peter Shirley, François Sillion, and Richard Szeliski for reviewing early versions of the entire monograph and providing many helpful comments.

Eric Stollnitz
Tony DeRose
David Salesin

Seattle, Washington
December 1995