



**Applied
Mathematical
Sciences**

153

**Stanley Osher
Ronald Fedkiw**

Level Set Methods and Dynamic Implicit Surfaces



Springer

Stanley Osher Ronald Fedkiw

Level Set Methods and Dynamic Implicit Surfaces

With 99 Figures, Including 24 in Full Color



Springer

Stanley Osher
Department of Mathematics
University of California
at Los Angeles
Los Angeles, CA 90095-1555
USA
sjo@math.ucla.edu

Ronald Fedkiw
Department of Computer Science
Stanford University
Stanford, CA 94305-9020
USA
fedkiw@cs.stanford.edu

Editors:

S.S. Antman
Department of Mathematics
and
Institute for Physical Science
and Technology
University of Maryland
College Park, MD 20742-4015
USA
ssa@math.umd.edu

J.E. Marsden
Control and Dynamical
Systems, 107-81
California Institute of
Technology
Pasadena, CA 91125
USA
marsden@cds.caltech.edu

L. Sirovich
Division of Applied
Mathematics
Brown University
Providence, RI 02912
USA
chico@camelot.mssm.edu

Cover photos: Top left and right, hand and rat brain — Duc Nguyen and Hong-Kai Zhao. Center campfire — Duc Nguyen and Nick Rasmussen and Industrial Light and Magic. Lower left and center, water glasses — Steve Marschner and Doug Enright.

Mathematics Subject Classification (2000): 65Mxx, 65C20, 65D17, 65-02, 65V10, 73V

Library of Congress Cataloging-in-Publication Data

Osher, Stanley.

Level set methods and dynamic implicit surfaces / Stanley Osher, Ronald Fedkiw
p. cm. — (Applied mathematical sciences ; 153)

Includes bibliographical references and index.

ISBN 0-387-95482-1 (alk. paper)

1. Level set methods. 2. Implicit functions. I. Fedkiw, Ronald P., 1968– II. Title.
III. Applied mathematical sciences (Springer-Verlag New York Inc.) ; v. 153

QA1.A647 vol. 153

[QC173.4]

510s–dc21

[515'.8]

2002020939

ISBN 0-387-95482-1

Printed on acid-free paper.

© 2003 Springer-Verlag New York, Inc.

All rights reserved. This work may not be translated or copied in whole or in part without the written permission of the publisher (Springer-Verlag New York, Inc., 175 Fifth Avenue, New York, NY 10010, USA), except for brief excerpts in connection with reviews or scholarly analysis. Use in connection with any form of information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed is forbidden.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

Printed in the United States of America.

9 8 7 6 5 4 3 2

SPIN 10920466

www.springer-ny.com

Springer-Verlag New York Berlin Heidelberg

A member of BertelsmannSpringer Science+Business Media GmbH

Dedicated with love to Katy, Brittany, and Bobbie

Preface

Scope, Aims, and Audiences

This book, *Level Set Methods and Dynamic Implicit Surfaces* is designed to serve two purposes:

Parts I and II introduce the reader to implicit surfaces and level set methods. We have used these chapters to teach introductory courses on the material to students with little more than a fundamental math background. No prior knowledge of partial differential equations or numerical analysis is required. These first eight chapters include enough detailed information to allow students to create working level set codes from scratch.

Parts III and IV of this book are based on a series of papers published by us and our colleagues. For the sake of brevity, a few details have been occasionally omitted. These chapters do include thorough explanations and enough of the significant details along with the appropriate references to allow the reader to get a firm grasp on the material.

This book is an introduction to the subject. We have given examples of the utility of the method to a diverse (but by no means complete) collection of application areas. We have also tried to give complete numerical recipes and a self-contained course in the appropriate numerical analysis. We believe that this book will enable users to apply the techniques presented here to real problems.

The level set method has been used in a rapidly growing number of areas, far too many to be represented here. These include epitaxial growth, optimal design, CAD, MEMS, optimal control, and others where the simulation

of moving interfaces plays a key role in the problem to be solved. A search of “level set methods” on the Google website (which gave over 2,700 responses as of May 2002) will give an interested reader some idea of the scope and utility of the method. In addition, some exciting advances in the technology have been made since we began writing this book. We hope to cover many of these topics in a future edition. In the meantime you can find some exciting animations and moving images as well as links to more relevant research papers via our personal web sites: <http://graphics.stanford.edu/~fedkiw> and <http://www.math.ucla.edu/~sjo/>.

Acknowledgments

Many people have helped us in this effort. We thank the following colleagues in particular: Steve Marschner, Paul Romburg, Gary Hower, and Steve Ruuth for proofreading parts of the manuscript, Peter Smereka and Li-Tien Cheng for providing figures for the chapter on Codimension-Two Objects, Myungjoo Kang for providing figures for the chapter on Motion Involving Mean Curvature and Motion in the Normal Direction, Antonio Marquina and Frederic Gibou for help with the chapter on Image Restoration, Hong-Kai Zhao for help with chapter 13, Reconstruction of Surfaces from Unorganized Data Points, and Luminita Vese for help with the chapter on Snakes, Active Contours, and Segmentation. We particularly thank Barry Merriman for his extremely valuable collaboration on much of the research described here. Of course we have benefitted immensely from collaborations and discussions with far too many people to mention. We hope these colleagues and friends forgive us for omitting their names.

We would like to thank the following agencies for their support during this period: ONR, AFOSR, NSF, ARO, and DARPA. We are particularly grateful to Dr. Wen Masters of ONR for suggesting and believing in this project and for all of her encouragement during some of the more difficult times.

Finally, we thank our families and friends for putting up with us during this exciting, but stressful period.

Los Angeles, California
Stanford, California

Stanley Osher
Ronald Fedkiw

Contents

Preface	vii
Color Insert	(facing page 146)
I Implicit Surfaces	1
1 Implicit Functions	3
1.1 Points	3
1.2 Curves	4
1.3 Surfaces	7
1.4 Geometry Toolbox	8
1.5 Calculus Toolbox	13
2 Signed Distance Functions	17
2.1 Introduction	17
2.2 Distance Functions	17
2.3 Signed Distance Functions	18
2.4 Examples	19
2.5 Geometry and Calculus Toolboxes	21
II Level Set Methods	23
3 Motion in an Externally Generated Velocity Field	25
3.1 Convection	25

3.2	Upwind Differencing	29
3.3	Hamilton-Jacobi ENO	31
3.4	Hamilton-Jacobi WENO	33
3.5	TVD Runge-Kutta	37
4	Motion Involving Mean Curvature	41
4.1	Equation of Motion	41
4.2	Numerical Discretization	44
4.3	Convection-Diffusion Equations	45
5	Hamilton-Jacobi Equations	47
5.1	Introduction	47
5.2	Connection with Conservation Laws	48
5.3	Numerical Discretization	49
5.3.1	Lax-Friedrichs Schemes	50
5.3.2	The Roe-Fix Scheme	52
5.3.3	Godunov's Scheme	54
6	Motion in the Normal Direction	55
6.1	The Basic Equation	55
6.2	Numerical Discretization	57
6.3	Adding a Curvature-Dependent Term	59
6.4	Adding an External Velocity Field	59
7	Constructing Signed Distance Functions	63
7.1	Introduction	63
7.2	Reinitialization	64
7.3	Crossing Times	65
7.4	The Reinitialization Equation	65
7.5	The Fast Marching Method	69
8	Extrapolation in the Normal Direction	75
8.1	One-Way Extrapolation	75
8.2	Two-Way Extrapolation	76
8.3	Fast Marching Method	76
9	Particle Level Set Method	79
9.1	Eulerian Versus Lagrangian Representations	79
9.2	Using Particles to Preserve Characteristics	82
10	Codimension-Two Objects	87
10.1	Intersecting Two Level Set Functions	87
10.2	Modeling Curves in \mathbb{R}^3	87
10.3	Open Curves and Surfaces	90
10.4	Geometric Optics in a Phase-Space-Based Level Set Framework	90

III Image Processing and Computer Vision 95

11 Image Restoration 97

- 11.1 Introduction to PDE-Based Image Restoration 97
- 11.2 Total Variation-Based Image Restoration 99
- 11.3 Numerical Implementation of TV Restoration 103

12 Snakes, Active Contours, and Segmentation 119

- 12.1 Introduction and Classical Active Contours 119
- 12.2 Active Contours Without Edges 121
- 12.3 Results 124
- 12.4 Extensions 124

13 Reconstruction of Surfaces from Unorganized Data Points 139

- 13.1 Introduction 139
- 13.2 The Basic Model 140
- 13.3 The Convection Model 142
- 13.4 Numerical Implementation 142

IV Computational Physics 147

14 Hyperbolic Conservation Laws and Compressible Flow 149

- 14.1 Hyperbolic Conservation Laws 149
 - 14.1.1 Bulk Convection and Waves 150
 - 14.1.2 Contact Discontinuities 151
 - 14.1.3 Shock Waves 152
 - 14.1.4 Rarefaction Waves 153
- 14.2 Discrete Conservation Form 154
- 14.3 ENO for Conservation Laws 155
 - 14.3.1 Motivation 155
 - 14.3.2 Constructing the Numerical Flux Function 157
 - 14.3.3 ENO-Roe Discretization (Third-Order Accurate) 158
 - 14.3.4 ENO-LLF Discretization (and the Entropy Fix) 159
- 14.4 Multiple Spatial Dimensions 160
- 14.5 Systems of Conservation Laws 160
 - 14.5.1 The Eigensystem 161
 - 14.5.2 Discretization 162
- 14.6 Compressible Flow Equations 163
 - 14.6.1 Ideal Gas Equation of State 164
 - 14.6.2 Eigensystem 164
 - 14.6.3 Numerical Approach 165

15 Two-Phase Compressible Flow	167
15.1 Introduction	167
15.2 Errors at Discontinuities	168
15.3 Rankine-Hugoniot Jump Conditions	169
15.4 Nonconservative Numerical Methods	171
15.5 Capturing Conservation	172
15.6 A Degree of Freedom	172
15.7 Isobaric Fix	173
15.8 Ghost Fluid Method	175
15.9 A Robust Alternative Interpolation	183
16 Shocks, Detonations, and Deflagrations	189
16.1 Introduction	189
16.2 Computing the Velocity of the Discontinuity	190
16.3 Limitations of the Level Set Representation	191
16.4 Shock Waves	191
16.5 Detonation Waves	193
16.6 Deflagration Waves	195
16.7 Multiple Spatial Dimensions	196
17 Solid-Fluid Coupling	201
17.1 Introduction	201
17.2 Lagrange Equations	203
17.3 Treating the Interface	204
18 Incompressible Flow	209
18.1 Equations	209
18.2 MAC Grid	210
18.3 Projection Method	212
18.4 Poisson Equation	213
18.5 Simulating Smoke for Computer Graphics	214
19 Free Surfaces	217
19.1 Description of the Model	217
19.2 Simulating Water for Computer Graphics	218
20 Liquid-Gas Interactions	223
20.1 Modeling	223
20.2 Treating the Interface	224
21 Two-Phase Incompressible Flow	227
21.1 Introduction	227
21.2 Jump Conditions	230
21.3 Viscous Terms	232
21.4 Poisson Equation	235

22 Low-Speed Flames	239
22.1 Reacting Interfaces	239
22.2 Governing Equations	240
22.3 Treating the Jump Conditions	241
23 Heat Flow	249
23.1 Heat Equation	249
23.2 Irregular Domains	250
23.3 Poisson Equation	251
23.4 Stefan Problems	254
References	259
Index	271

Part I

Implicit Surfaces

In the next two chapters we introduce implicit surfaces and illustrate a number of useful properties, focusing on those that will be of use to us later in the text. A good general review can be found in [16]. In the first chapter we discuss those properties that are true for a general implicit representation. In the second chapter we introduce the notion of a signed distance function with a Euclidean distance metric and a “ \pm ” sign used to indicate the inside and outside of the surface.

1

Implicit Functions

1.1 Points

In one spatial dimension, suppose we divide the real line into three distinct pieces using the points $x = -1$ and $x = 1$. That is, we define $(-\infty, -1)$, $(-1, 1)$, and $(1, \infty)$ as three separate subdomains of interest, although we regard the first and third as two disjoint pieces of the same region. We refer to $\Omega^- = (-1, 1)$ as the *inside* portion of the domain and $\Omega^+ = (-\infty, -1) \cup (1, \infty)$ as the *outside* portion of the domain. The border between the inside and the outside consists of the two points $\partial\Omega = \{-1, 1\}$ and is called the *interface*. In one spatial dimension, the inside and outside regions are one-dimensional objects, while the interface is less than one-dimensional. In fact, the points making up the interface are zero-dimensional. More generally, in \mathbb{R}^n , subdomains are n -dimensional, while the interface has dimension $n - 1$. We say that the interface has codimension one.

In an *explicit* interface representation one explicitly writes down the points that belong to the interface as we did above when defining $\partial\Omega = \{-1, 1\}$. Alternatively, an *implicit* interface representation defines the interface as the isocontour of some function. For example, the zero isocontour of $\phi(x) = x^2 - 1$ is the set of all points where $\phi(x) = 0$; i.e., it is exactly $\partial\Omega = \{-1, 1\}$. This is shown in Figure 1.1. Note that the implicit function $\phi(x)$ is defined throughout the one-dimensional domain, while the isocontour defining the interface is one dimension lower. More generally, in \mathbb{R}^n , the implicit function $\phi(\vec{x})$ is defined on all $\vec{x} \in \mathbb{R}^n$, and its isocontour has dimension $n - 1$. Initially, the implicit representation might seem




Figure 1.1. Implicit function $\phi(x) = x^2 - 1$ defining the regions Ω^- and Ω^+ as well as the boundary $\partial\Omega$

wasteful, since the implicit function $\phi(\vec{x})$ is defined on all of \Re^n , while the interface has only dimension $n - 1$. However, we will see that a number of very powerful tools are readily available when we use this representation.

Above, we chose the $\phi(x) = 0$ isocontour to represent the lower-dimensional interface, but there is nothing special about the zero isocontour. For example, the $\hat{\phi}(x) = 1$ isocontour of $\hat{\phi}(x) = x^2$, defines the same interface, $\partial\Omega = \{-1, 1\}$. In general, for any function $\hat{\phi}(\vec{x})$ and an arbitrary isocontour $\hat{\phi}(\vec{x}) = a$ for some scalar $a \in \Re$, we can define $\phi(\vec{x}) = \hat{\phi}(\vec{x}) - a$, so that the $\phi(\vec{x}) = 0$ isocontour of ϕ is identical to the $\hat{\phi}(\vec{x}) = a$ isocontour of $\hat{\phi}$. In addition, the functions ϕ and $\hat{\phi}$ have identical properties up to a scalar translation a . Moreover, the partial derivatives of ϕ are the same as the partial derivatives of $\hat{\phi}$, since the scalar vanishes upon differentiation. Thus, throughout the text all of our implicit functions $\phi(\vec{x})$ will be defined so that the $\phi(\vec{x}) = 0$ isocontour represents the interface (unless otherwise specified).

1.2 Curves

In two spatial dimensions, our lower-dimensional interface is a curve that separates \Re^2 into separate subdomains with nonzero areas. Here we are limiting our interface curves to those that are *closed*, so that they have clearly defined interior and exterior regions. As an example, consider $\phi(\vec{x}) =$




Figure 1.2. Implicit representation of the curve $x^2 + y^2 = 1$.

$x^2 + y^2 - 1$, where the interface defined by the $\phi(\vec{x}) = 0$ isocontour is the unit circle defined by $\partial\Omega = \{\vec{x} \mid |\vec{x}| = 1\}$. The interior region is the unit open disk $\Omega^- = \{\vec{x} \mid |\vec{x}| < 1\}$, and the exterior region is $\Omega^+ = \{\vec{x} \mid |\vec{x}| > 1\}$. These regions are depicted in Figure 1.2. The explicit representation of this interface is simply the unit circle defined by $\partial\Omega = \{\vec{x} \mid |\vec{x}| = 1\}$.

In two spatial dimensions, the explicit interface definition needs to specify all the points on a curve. While in this case it is easy to do, it can be somewhat more difficult for general curves. In general, one needs to parameterize the curve with a vector function $\vec{x}(s)$, where the parameter s is in $[s_o, s_f]$. The condition that the curve be closed implies that $\vec{x}(s_o) = \vec{x}(s_f)$.

While it is convenient to use analytical descriptions as we have done so far, complicated two-dimensional curves do not generally have such simple representations. A convenient way of approximating an explicit representation is to discretize the parameter s into a finite set of points $s_o < \dots < s_{i-1} < s_i < s_{i+1} < \dots < s_f$, where the subintervals $[s_i, s_{i+1}]$ are not necessarily of equal size. For each point s_i in parameter space, we then store the corresponding two-dimensional location of the curve denoted by $\vec{x}(s_i)$. As the number of points in the discretized parameter space is increased, so is the resolution (detail) of the two-dimensional curve.

The implicit representation can be stored with a discretization as well, except now one needs to discretize all of \mathbb{R}^2 , which is impractical, since it is unbounded. Instead, we discretize a bounded subdomain $D \subset \mathbb{R}^2$. Within this domain, we choose a finite set of points (x_i, y_i) for $i = 1, \dots, N$ to discretely approximate the implicit function ϕ . This illustrates a drawback of the implicit surface representation. Instead of resolving a one-dimensional

interval $[s_o, s_f]$, one needs to resolve a two-dimensional region D . More generally, in \mathbb{R}^n , a discretization of an explicit representation needs to resolve only an $(n - 1)$ -dimensional set, while a discretization of an implicit representation needs to resolve an n -dimensional set. This can be avoided, in part, by placing all the points \vec{x} very close to the interface, leaving the rest of D unresolved. Since only the $\phi(\vec{x}) = 0$ isocontour is important, only the points \vec{x} near this isocontour are actually needed to accurately represent the interface. The rest of D is unimportant. Clustering points near the interface is a *local* approach to discretizing implicit representations. (We will give more details about local approaches later.) Once we have chosen the set of points that make up our discretization, we store the values of the implicit function $\phi(\vec{x})$ at each of these points.

Neither the explicit nor the implicit discretization tells us where the interface is located. Instead, they both give information at sample locations. In the explicit representation, we know the location of a finite set of points on the curve, but do not know the location of the remaining infinite set of points (on the curve). Usually, interpolation is used to approximate the location of points not represented in the discretization. For example, piecewise polynomial interpolation can be used to determine the shape of the interface between the data points. Splines are usually appropriate for this. Similarly, in the implicit representation we know the values of the implicit function ϕ at only a finite number of points and need to use interpolation to find the values of ϕ elsewhere. Even worse, here we may not know the location of any of the points on the interface, unless we have luckily chosen data points \vec{x} where $\phi(\vec{x})$ is exactly equal to zero. In order to locate the interface, the $\phi(\vec{x}) = 0$ isocontour needs to be interpolated from the known values of ϕ at the data points. This is a rather standard procedure accomplished by a variety of contour plotting routines.

The set of data points where the implicit function ϕ is defined is called a *grid*. There are many ways of choosing the points in a grid, and these lead to a number of different types of grids, e.g., unstructured, adaptive, curvilinear. By far, the most popular grids, are Cartesian grids defined as $\{(x_i, y_j) \mid 1 \leq i \leq m, 1 \leq j \leq n\}$. The natural orderings of the x_i and y_j are usually used for convenience. That is, $x_1 < \cdots < x_{i-1} < x_i < x_{i+1} < \cdots < x_m$ and $y_1 < \cdots < y_{j-1} < y_j < y_{j+1} < \cdots < y_n$. In a *uniform* Cartesian grid, all the subintervals $[x_i, x_{i+1}]$ are equal in size, and we set $\Delta x = x_{i+1} - x_i$. Likewise, all the subintervals $[y_j, y_{j+1}]$ are equal in size, and we set $\Delta y = y_{j+1} - y_j$. Furthermore, it is usually convenient to choose $\Delta x = \Delta y$ so that the approximation errors are the same in the x -direction as they are in the y -direction. By definition, Cartesian grids imply the use of a rectangular domain $D = [x_1, x_m] \times [y_1, y_n]$. Again, since ϕ is important only near the interface, a local approach would indicate that many of the grid points are not needed, and the implicit representation can be optimized by storing only a subset of a uniform Cartesian grid. The Cartesian grid points that are not sufficiently near the interface can be discarded.