

Foundations of 3D Computer Graphics

— -1
— 0
— 1

— -1
— 0
— 1

FOUNDATIONS OF 3D COMPUTER GRAPHICS

Steven J. Gortler

**The MIT Press
Cambridge, Massachusetts
London, England**

—1
— 0
— 1

© 2012 Steven J. Gortler

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

MIT Press books may be purchased at special quantity discounts for business or sales promotional use. For information, please email special_sales@mitpress.mit.edu or write to Special Sales Department, The MIT Press, 55 Hayward Street, Cambridge, MA 02142.

This book was set in Syntax and Times Roman by Westchester Book Group. Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

10 9 8 7 6 5 4 3 2 1

—1
— 0
— 1

To A.L.G., F.J.G, and O.S.G.

— -1
— 0
— 1

— -1
— 0
— 1

Contents

Preface xiii

I GETTING STARTED 1

1 Introduction 3

1.1 OpenGL 3
Exercises 8

2 Linear 9

2.1 Geometric Data Types 9
2.2 Vectors, Coordinate Vectors, and Bases 11
2.3 Linear Transformations and 3 by 3 Matrices 12
2.4 Extra Structure 15
2.5 Rotations 17
2.6 Scales 19
Exercises 20

3 Affine 21

3.1 Points and Frames 21
3.2 Affine Transformations and 4 by 4 Matrices 22
3.3 Applying Linear Transformations to Points 24
3.4 Translations 25
3.5 Putting Them Together 25
3.6 Normals 26
Exercise 28

4 Respect 29

4.1 The Frame Is Important 29
4.2 Multiple Transformations 31
Exercises 33

5 Frames in Graphics 35

5.1 World, Object, and Eye Frames 35
5.2 Moving Things Around 36

—1
— 0
— 1

5.3	Scales	39
5.4	Hierarchy	40
	Exercises	43
6	Hello World 3D	45
6.1	Coordinates and Matrices	45
6.2	Drawing a Shape	46
6.3	The Vertex Shader	51
6.4	What Happens Next	52
6.5	Placing and Moving with Matrices	53
	Exercises	54
II	ROTATIONS AND INTERPOLATION	57
7	Quaternions (a Bit Technical)	59
7.1	Interpolation	59
7.2	The Representation	63
7.3	Operations	64
7.4	Power	65
7.5	Code	68
7.6	Putting Back the Translations	68
	Exercises	72
8	Balls: Track and Arc	73
8.1	The Interfaces	73
8.2	Properties	74
8.3	Implementation	75
	Exercise	75
9	Smooth Interpolation	77
9.1	Cubic Bezier Functions	77
9.2	Catmull–Rom Splines	80
9.3	Quaternion Splining	82
9.4	Other Splines	82
9.5	Curves in Space	83
	Exercises	85
III	CAMERAS AND RASTERIZATION	87
10	Projection	89
10.1	Pinhole Camera	89
10.2	Basic Mathematical Model	91
10.3	Variations	92
10.4	Context	98
	Exercises	98

— 1
— 0
— 1

- 11 Depth 101**
 - 11.1 Visibility 101
 - 11.2 Basic Mathematical Model 102
 - 11.3 Near and Far 105
 - 11.4 Code 107
 - Exercises 108

- 12 From Vertex to Pixel 109**
 - 12.1 Clipping 109
 - 12.2 Backface Culling 113
 - 12.3 Viewport 114
 - 12.4 Rasterization 116
 - Exercises 118

- 13 Varying Variables (Tricky) 119**
 - 13.1 Motivating the Problem 119
 - 13.2 Rational Linear Interpolation 121
 - Exercises 123

- IV PIXELS AND SUCH 125**

- 14 Materials 127**
 - 14.1 Basic Assumptions 127
 - 14.2 Diffuse 130
 - 14.3 Shiny 131
 - 14.4 Anisotropy 133
 - Exercise 136

- 15 Texture Mapping 137**
 - 15.1 Basic Texturing 137
 - 15.2 Normal Mapping 139
 - 15.3 Environment Cube Maps 140
 - 15.4 Projector Texture Mapping 141
 - 15.5 Multipass 144
 - Exercise 148

- 16 Sampling 149**
 - 16.1 Two Models 149
 - 16.2 The Problem 150
 - 16.3 The Solution 150
 - 16.4 Alpha 155
 - Exercise 160

- 17 Reconstruction 161**
 - 17.1 Constant 161
 - 17.2 Bilinear 162

____-1
 ____ 0
 ____ 1

17.3	Basis Functions	163
	Exercises	165
18	Resampling	167
18.1	Ideal Resampling	167
18.2	Blow Up	169
18.3	Mip Map	169
V	ADVANCED TOPICS	173
19	Color	175
19.1	Simple Biophysical Model	175
19.2	Mathematical Model	180
19.3	Color Matching	183
19.4	Bases	185
19.5	Reflection Modeling	187
19.6	Adaptation	190
19.7	Nonlinear Color	191
	Exercises	195
20	What Is Ray Tracing?	197
20.1	Loop Ordering	197
20.2	Intersection	199
20.3	Secondary Rays	201
	Exercises	202
21	Light (Technical)	205
21.1	Units	205
21.2	Reflection	210
21.3	Light Simulation	214
21.4	Sensors	221
21.5	Integration Algorithms	222
21.6	More General Effects	224
	Exercises	225
22	Geometric Modeling: Basic Intro	227
22.1	Triangle Soup	227
22.2	Meshes	227
22.3	Implicit Surfaces	229
22.4	Volume	231
22.5	Parametric Patches	231
22.6	Subdivision Surfaces	232
	Exercises	237

23 Animation: Not Even an Introduction 239

- 23.1 Interpolation 239
- 23.2 Simulation 241
- 23.3 Human Locomotion 246
 - Exercise 246

APPENDIXES 247**A Hello World 2D 249**

- A.1 APIs 249
- A.2 Main Program 250
- A.3 Adding Some Interaction 257
- A.4 Adding a Texture 259
- A.5 What's Next 261
 - Exercises 262

B Affine Functions 263

- B.1 2D Affine 263
- B.2 3D Affine 264
- B.3 Going Up 264
- B.4 Going Down 265
- B.5 Going Sideways 265
 - Exercises 266

References 267

Index 000

— -1
— 0
— 1

Preface

This book developed out of an introductory computer graphics course I have been teaching at Harvard since 1996. Over the years, I have had the pleasure of teaching many amazing students. During class, these students have asked many good questions. In light of these questions, I often realized that some of my explanations in class were a bit sloppy and that I didn't fully understand the topic I had just tried to explain. This would often lead me to rethink the material and change the way I taught it the next time around. Many of these ideas have found their way into this book. Throughout the course of the book, I cover mostly standard material but with an emphasis on understanding some of the more subtle concepts involved.

In this book, we will introduce the basic algorithmic technology needed to produce three-dimensional (3D) computer graphics. We will cover the basic ideas of how 3D shapes are represented and moved around algorithmically. We will cover how a camera can be algorithmically modeled turning this 3D data into a two-dimensional (2D) image made up of a discrete set of dots, or *pixels*, on a screen. Later in the book, we will cover some advanced topics on the basics of color and light representations. We will also briefly introduce some advanced topics on light simulation for producing photo-realistic images, on various ways of dealing with geometric representations, and on producing animated computer graphics.

In this book, we include material that is both above and below the API-hood. Much of the material (especially early on) is stuff you simply need to know to do 3D computer graphics. But we also spend time to explain what is going on inside of OpenGL. This is necessary to understand in order to be a highly competent computer graphics programmer. But also, it is simply fascinating to learn the hows and whys of our amazing computer graphics computational infrastructure.

We will not cover the hardware and compiler aspects of computer graphics in this book. We will also not focus on 2D computer graphics or human-computer interfaces. These topics are all interesting in their own rights but are fairly distinct from the algorithmic side of 3D computer graphics.

In this book, we structure our explanations around OpenGL, a real-time “rasterization-based” rendering environment. We have done this (rather than, say, a “ray tracing-based”

—1
— 0
— 1

environment) because so much of computer graphics is done in this setting. Anyone, for example, who works in 3D video games needs to master this material (and more). We have chosen the OpenGL API (with the GLSL shading language) in particular, as it can be run on a wide variety of computing platforms.

This book is intended for upper-level computer science/math/physics undergraduate students with at least a year of programming under their belts and at least a rudimentary understanding of linear algebra.

For the Professor

In the following paragraphs, I will describe some of the subtle issues that require some care to get right and are often not taught clearly. I hope that students will master these topics from this book.

Chapters 2–4: In computer graphics, we need to think about points and vectors from both a coordinate-free and a coordinate-full approach. We need to use coordinates to obtain a concrete representation ultimately resulting in our rendered images. But it is often important to represent and transform our points with respect to different coordinate systems. As such it is important to

- Distinguish between a geometric point and the coordinates used to represent that point with respect to some frame.
- Use a notation to explicitly keep track of the basis used for a set of coordinates.
- Distinguish in our notation between matrix equations that represent basis changes and matrix expressions that represent geometric transformations being applied to points.

This ultimately leads to what we call the *left-of* rule, which allows us to interpret matrix expressions and understand with respect to which basis a transformation is acting.

It is our hope that by mastering this explicit notational framework, a student can easily figure out how to do complicated transformations. This is in contrast to the “try lots of orderings and inverses of sequences of matrices until the program does the right thing” approach. One loose inspiration for our approach is the manuscript by Tony DeRose [16].

Chapters 5 and 6: We describe an organized framework for dealing with frames in computer graphics and how this translates into simple 3D OpenGL code. In particular, we derive useful ways to move objects around using a “mixed auxiliary frame.” This allows us to, say, rotate an object correctly about its own center but in directions that correspond naturally to those on the screen.

Chapter 7: This is a simple and straightforward description of the quaternion representation for rotations. We also derive how properly to combine quaternions and translation vectors to define a rigid-body transformation data type that multiplies and inverts just like matrices.

____—1
 ____ 0
 ____ 1

Chapter 8: This is a simple and straightforward description of the trackball and arcball rotation interface. We also show why the trackball interface is mouse path-invariant.

Chapter 9: We do a quick and dirty introduction to Bezier and Catmull-Rom splines.

Chapters 10–13: In these chapters, we describe how camera projection is modeled using 4 by 4 matrices. We also describe the fixed function operations in OpenGL. We pay special attention to deriving the correct formulas for interpolating the varying variables. Some of the background on affine functions and interpolation is relegated to appendix B. Many of the subtleties here are described nicely in essays by Jim Blinn [5]. In these chapters, we do not cover details about clipping or rasterization algorithms.

Chapters 14 and 15: We give some simple example shaders for diffuse, shiny, and anisotropic materials. We also point to more advanced real-time rendering techniques such as multipass rendering and shadow mapping. These sections are admittedly too brief. Presumably, students pursuing more aggressive rendering projects will need to learn a lot more about the ever-evolving techniques (and tricks) used in modern real-time rendering.

Chapters 16–18: We cover the basics of why filtering is needed to compute discrete images and how basis functions are used to reconstruct continuous ones. In particular, we show how these two concepts need to be combined to do filtering properly during texture mapping (à la Paul Heckbert's M.S. thesis [28]). We do not delve into the details of Fourier analysis here, as we think it would pull us a bit too far off the main narrative (and, in the end, we use box filters anyway).

Chapter 19: We describe the basic theory of human color perception. From a mathematical perspective, we attempt to be clear about the very definition of what a color is and why such things form a linear space. A related treatment to ours can be found in Feynman's lectures [20]. For many of the technical issues of color computations, we rely on the color FAQ of Charles Poynton [58].

Chapter 20: For completeness, we briefly describe ray tracing computations. As this is not the focus of the course, we only touch on the topic.

Chapter 21: As an introduction to advanced topics in photo-realistic rendering, we do a careful introduction to the physical units for describing light and to the reflection, rendering, and measurement equations. One thing we pay close attention to, for example, is why reflection is measured in the particular units used. A good reference on these basics and more is Eric Veach's Ph.D. thesis [71].

Chapter 22: We outline some of the ways that surfaces are modeled and represented in computer graphics. This is a nontechnical discussion that gives a quick introduction to this rich topic. We do go into enough detail to be able to implement Catmull-Rom subdivision surfaces (assuming you have a mesh data structure handy), as these are a quick and easy way to represent a broad family of surfaces.

Chapter 23: We outline some of the ways animation is done in computer graphics. Again, this is a nontechnical discussion that gives a quick introduction to this rich topic.

____—1
 ____ 0
 ____ 1

One nice place to start for this material is Adrien Treuille’s course on the CMU website (<http://www.cs.cmu.edu/~15869-f10>).

Appendix A: We try to get up and running as painlessly as possible with a first OpenGL program. Because we will be using a modern version of OpenGL, there is thankfully not much API left to teach anymore. All OpenGL needs to do is manage the shader programs, vertex buffer objects, and textures. We do not teach any of the deprecated elements of old OpenGL, as they are not used in modern computer graphics programming. Appendix A may be read at any time before chapter 6.

Appendix B: We summarize important facts about affine functions. This is helpful for reading chapter 12 and essential for reading chapter 13.

This book covers a bit more than what would be possible to do in a one-semester course. It does not attempt to be an encyclopedic reference to all of the rich topics in computer graphics theory and practice. Additionally, this book stays close to material in wide use today. We do not cover the many current research results and ideas that may one day become part of standard practice.

The website for the book may be found at <http://mitpress.mit.edu/foundations>.

Acknowledgments

During the development of this book over the past year, I received lots of helpful input from Hamilton Chong, Guillermo Diez-Canas, and Dylan Thurston.

Helpful advice also came from Julie Dorsey, Hugues Hoppe, Zoe Wood, Yuanchen Zhu, and Todd Zickler.

Other comments and assistance came from Gilbert Bernstein, Fredo Durand, Ladislav Kavan, Michael Kazhdan, and Peter-Pike Sloan.

During the development of the course, I was greatly assisted by Xiangfeng Gu, Danil Kirsanov, Chris Mihelich, Pedro Sander, and Abraham Stone.

Over the years, the course has had many talented students acting as teaching assistants. They have contributed in many ways to the evolution of this material and include Brad Andalman, Keith Baldwin, Forrester Cole, Ashley Eden, David Holland, Brad Kittenbrink, Sanjay Mavinkurve, Jay Moorthi, Doug Nachand, Brian Osserman, Ed Park, David Ryu, Razvan Surdulescu, and Geetika Tewari.

Finally, I thank my parents for their years of love and support.

— 1
— 0
— 1