Book Contents

Chapter 1: Analytical Geometry 1
Olga Sorkine-Hornung and Daniel Cohen-Or

In the first chapter, we will familiarize ourselves with some basic geometric tools and see how we can put them to practical use to solve several geometric problems. Instead of describing the tools directly, we do it through an interesting discussion of two possible ways to approach the geometric problem at hand: we can employ our geometric intuition and use geometric reasoning, or we can directly formalize everything and employ our algebraic skills to write down and solve some equations. The discussion leads to a presentation of linear geometric elements (points, lines, planes), and the means to manipulate them in common geometric applications that we encounter, such as distances, transformations, projections and more.

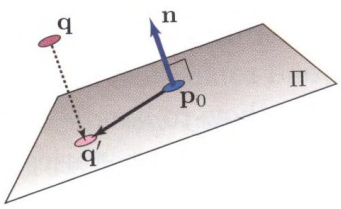

Chapter 2: Linear Algebra? 13
Daniel Cohen-Or, Olga Sorkine-Hornung and Chen Greif

In this chapter, we will review basic linear algebra notions that we learned in a basic linear algebra course, including vector spaces, orthogonal bases, subspaces, eigenvalues and eigenvectors. However, our main goal here is to convince the readers that these notions are really useful. Furthermore, we will see the close relation between linear algebra and geometry. The chapter will be driven by an important tool called singular value decomposition (SVD), to which we will devote a separate full chapter. To understand what an SVD is, we first need to understand the notions of bases, eigenvectors, and eigenvalues and to refresh some fundamentals of linear algebra with examples in geometric context.

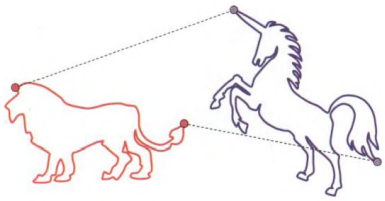

Chapter 3: Least-Squares Solutions 31
Niloy J. Mitra

When dealing with real-world data, simple patterns can often be submerged in noise and outliers. In this chapter, we will learn about basic data fitting using the least-squares method, first starting with simple line fitting before moving on to fitting low-order polynomials. Beyond robustness to noise, we will also learn how to handle outliers and look at basic robust statics.

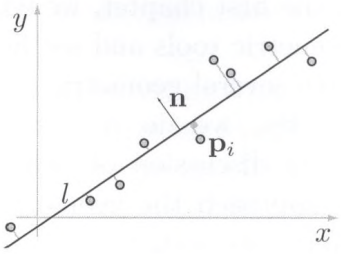

Chapter 4: PCA and SVD 47
Olga Sorkine-Hornung

In this chapter, we introduce two related tools from linear algebra that have become true workhorses in countless areas of science: principal component analysis (PCA) and singular value decomposition (SVD). These tools are extremely useful in geometric modeling, computer vision, image processing, computer graphics, machine learning and many other applications. We will see how to decompose a matrix into several factors that are easy to analyze and reveal important properties of the matrix and hence the data, or the problem in which the matrix arises. As in the whole book, the presentation is rather light, emphasizing the main principles without excessive rigor.

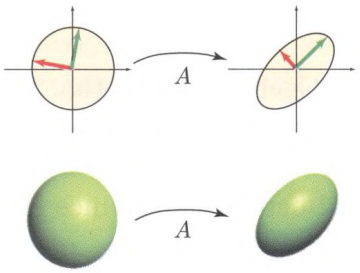

Chapter 5: Spectral Transform 63
Hao (Richard) Zhang

The use of signal transforms, such as the discrete Fourier or cosine transforms, is a classic topic in image and signal processing. In this chapter, we will learn how such transforms can be formulated and applied to the processing of 2D and 3D geometric shapes. The key concept to take away is the use of eigenvectors of discrete Laplacian operators as basis vectors to define spectral transforms for geometry. We will show how the Laplacian operators can be defined for 2D and 3D shapes, as well as a few applications of spectral transforms including geometry smoothing, enhancement and compression.

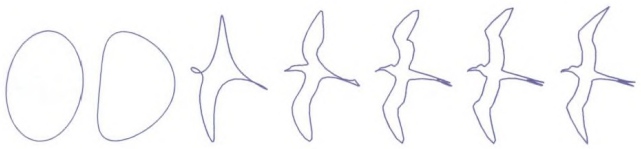

Chapter 6: Solution of Linear Systems 81
Chen Greif

In the solution of problems discussed in this book, a frequent task that arises is the need to solve a linear system. Understanding the properties of the matrix associated with the linear system is critical for guaranteeing speed and accuracy of the solution procedure. In this chapter, we provide an overview of linear system solvers. We describe direct methods and iterative methods, and discuss important criteria for the selection of a solution method, such as sparsity and positive definiteness. Important notions such as pivoting and preconditioning are explained, and a recipe is provided that helps in determining which solver should be used.

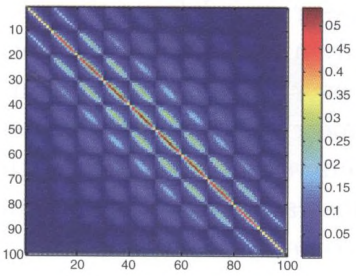

Chapter 7: Laplace and Poisson.....99
Daniel Cohen-Or and Gil Hoffer

In this chapter, we make use of the well-known equations of Laplace and Poisson. The two equations have an extremely simple form, and they are very useful in many diverse branches of mathematical physics. However, in this chapter, we will interpret them in the context of image processing. We will show some interesting image editing and geometric problems and how they can be solved by simple means using these equations.

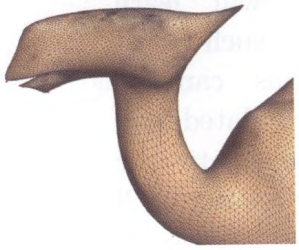

Chapter 8: Curvatures: A Differential Geometry Tool..117
Niloy J. Mitra and Daniel Cohen-Or

Local surface details, e.g., how "flat" a surface is locally, carry important information about the underlying object. Such information is critical for many applications in geometry processing, ranging from surface meshing, shape matching, surface reconstruction, scan alignment and detail-preserving deformation, to name only a few. In this chapter, we will cover the basics of differential geometry, particularly focusing on curvature estimates with some illustrative examples as an aid to geometry processing tasks.

Chapter 9: Dimensionality Reduction 131
Hao (Richard) Zhang and Daniel Cohen-Or

In this chapter, we will learn the concept, usefulness, and execution of dimensionality reduction. Generally speaking, we will seek to reduce the dimensionality of a given data set, mapping high-dimensional data into a lower-dimensional space to facilitate visualization, processing, or inference. We will present and discuss only a sample of dimensionality reduction techniques and illustrate them using visually intuitive examples, including face recognition, surface flattening and pose normalization of 3D shapes.

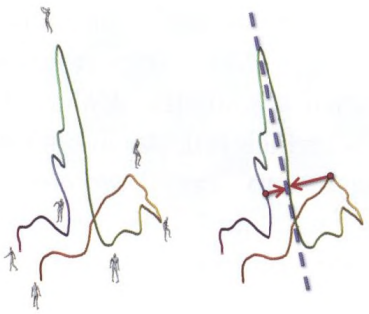

Chapter 10: Scattered Data Interpolation 147
Tao Ju

In this chapter, we visit the classical mathematical problem of obtaining a continuous function over a spatial domain from data at a few sample locations. The problem comes up in various geometric modeling scenarios, a good example of which is surface reconstruction. The chapter will eventually introduce the very useful radial basis functions (RBFs) as a smooth and efficient solution to the interpolation problem. However, to understand their usefulness, the chapter will go through a succession of methods with increasing sophistication, including piecewise linear interpolation and Shepherd's method.

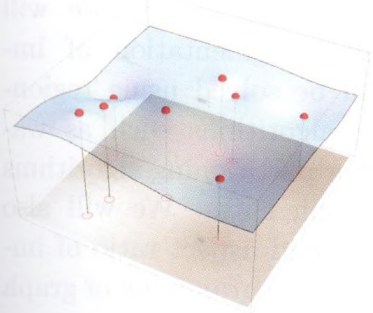

Chapter 11: Topology: How Are Objects Connected? 163
Niloy J. Mitra

In Chapter 8, we learned about local differential analysis of surfaces. In this chapter, we focus on global aspects. We will learn about what is meant by orientable surfaces or manifold surfaces. Most importantly, we will learn about the Euler characteristic, which links local curvature properties to global connectivity constraints, and comes up in a surprising range of applications.

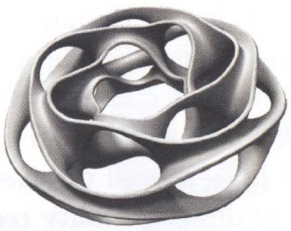

Chapter 12: Graphs and Images 177
Ariel Shamir

Graphs play an important role in many computer science fields and are also extensively used in imaging and graphics. This chapter concentrates on image processing and demonstrates how images can be represented by a graph. This allows translating problems of analysis and manipulation of images to well-known graph algorithms. Specifically, we will show how segmentation of images can be solved using region-growing algorithms such as watershed or partitioning algorithms using graph cuts. We will also show how intelligently changing the size and aspect ratio of images and video can be solved using dynamic programming or graph cuts.

Chapter 13: Skewing Scheme 205
Daniel Cohen-Or

In this chapter, we will show an example of the usefulness of number theory, or at least one of its known theorems. We will discuss mappings of numbers to a lattice, a problem that has practical applications in systems that require simultaneous, conflict-free access to elements distributed in different memory modules. Such mappings are also called *skewing schemes* since they skew the trivial mapping from element to memory. To understand these mappings, we will visit the notions of relatively prime numbers, and the greatest common divisor (gcd).

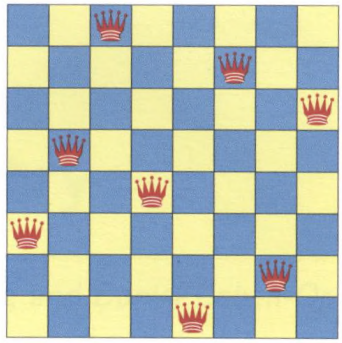

Bibliography 213

Index 223