1 Introduction

This chapter provides a brief summary of what geometric methods for signal and image analysis are all about. Several applications to imaging, computer graphics, and sensor networks are discussed and illustrated. The diversified nature of these applications is powerful testimony to the practical usage of geometric methods.

1.1 What is signal and image analysis?

Signal and image analysis refers to the extraction of meaningful information from signals and images using digital signal and image processing techniques. In signal processing, for example, we measure, manipulate, or analyze information about a signal [1]. Such a signal is defined as a function of one or more independent variables that carries information. Examples of signals are daily high temperatures measured over a month, voltages and currents in a circuit, stock prices, our voices, music and speech, images, videos, and emails. Image processing, on the other hand, is the study of any algorithm that takes an image as input and returns an image as output, such as image enhancement, segmentation, compression, inpainting, and feature detection, to name just a few [2]. Application domains of image processing abound and include medical imaging, biology, satellite imagery, and biometrics. Other areas closely related to signal and image processing are computer vision and computer graphics.

1.2 Why geometric methods?

In light of the successful use of signal and image processing principles and methodologies in a broad range of areas of exceptional social and economic value, there is a growing demand from within academia and industry to devise robust computational methodologies that uncover the key geometric and topological information from signals and images. Geometry is a branch of mathematics concerned with rigid form, size, and location of objects [3], whereas topology is one of the younger fields of mathematics concerned with properties that are preserved under continuous deformations of objects, such as deformations that involve stretching, shrinking, and twisting but no tearing or gluing [4]. In other words, topology is the study of the "shape" of curves and surfaces,

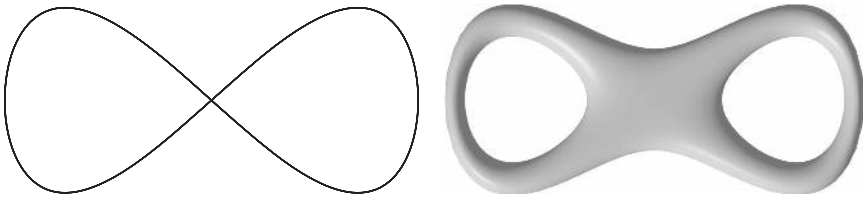

Figure 1.1 Curve (left); surface (right).

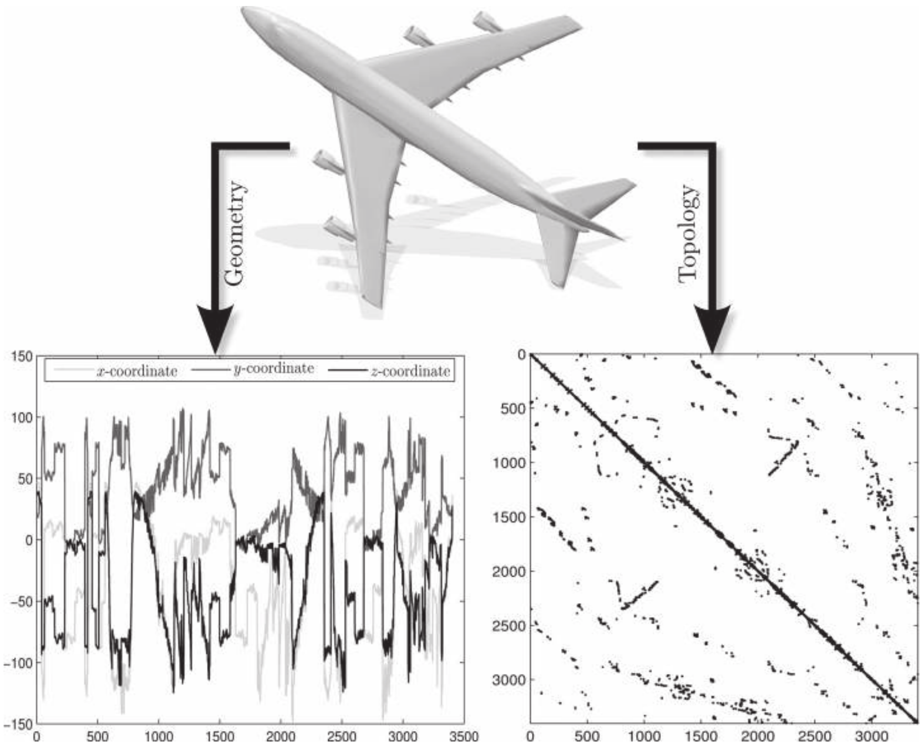

Figure 1.2 Geometric and topological information in a triangle mesh.

while geometry determines where, in a given coordinate system, each part is located. Figure 1.1 shows examples of a curve and surface.

In computer graphics, for example, a surface of a three-dimensional (3D) shape is usually modeled as a triangle mesh, which consists of a set of vertices and a set of triangles. The set of vertices encode geometry, while the set of triangles essentially encode topology (connectivity information), as Figure 1.2 illustrates. In the graphs of Figure 1.2, the one on the left displays the x-, y- and z-coordinate plots of the mesh vertices, while the one on the right depicts the adjacency matrix of the mesh. The adjacency matrix (also called connection matrix) is a sparse matrix that simply tells us whether two mesh vertices are connected or not.

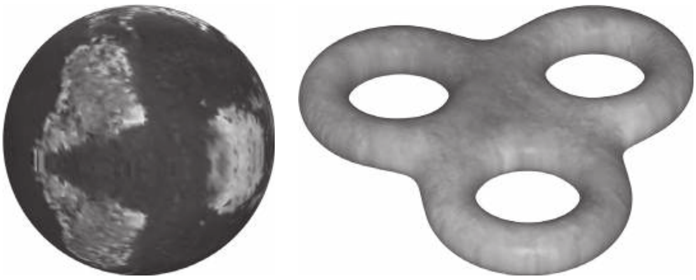

Figure 1.3 Examples of 2-manifolds: surface of Earth (left); surface of a pretzel (right).

In recent years, there has been a tremendous interest in developing computational geometric and topological methods for solving challenging signal and image analysis problems arising in a wide range of areas, including medicine, remote sensing, astronomy, robotics, security, and defense. This has been motivated, in large part, by the fact that virtually all objects contain geometric and topological information. However, our awareness of the importance of this information has been an evolutionary development. In this book, we aim at bringing these recent developments to the fore by providing a balanced coverage of both theoretical and practical issues. These computational geometric and topological methods combine classical mathematical techniques from geometry and algebraic topology with more recent algorithmic tools from data structure design and geometric computing. We will be especially interested in the discovery and application of geometric and topological properties of 3D surfaces to image analysis and geometry processing. Surfaces (also called 2-manifolds) are objects that are locally two-dimensional, i.e. around any point on the surface there is a neighborhood that locally looks like the plane. Examples of 2-manifolds include the surface of Earth and the surface of a pretzel, as shown in Figure 1.3.

The field of geometric computing for image analysis is concerned with the design, analysis, and implementation of algorithms for geometric problems that arise in a wide range of image processing applications [5]. Since the mid 1980s, geometric computing has surged as an independent field, with its own international conferences, workshops, and journals. In the early years, mostly theoretical foundations of geometric approaches were laid and fundamental research remains an important issue in the field. Meanwhile, as the field matured, applications and implementations of geometric methods in image analysis have increasingly received much attention and are gaining steam very rapidly.

Computational topology, on the other hand, is an emerging field of study that is devoted to the design of efficient algorithms for solving topological problems [6, 7], especially topological invariants that are of paramount importance to image analysis and geometry processing. Since the mid 1990s, computational topology has emerged as an independent field. Today, the fast-growing field of geometric and topological computing enjoys a broad and solid foundation upon which its future can be securely built. This important research stream offers great opportunities for researchers and practitioners to employ and integrate theories from geometry and topology, and draw on established theoretical and algorithmic frameworks. The present book is, in part, a response to these trends, offering a comprehensive reference which should be particularly useful for advanced undergraduates, for postgraduates interested in signal and image processing, and for practitioners in these fields.

1.3 Applications

Geometric and topological methods are widely used in a variety of applications including, but not limited to, signal and image analysis, surface denoising and compression, shape recognition, medical imaging, multimedia security, and sensor networks. This section provides a brief overview of some of these applications.

1.3.1 Image edge detection

Edges are prominent features in images and their analysis and detection are an essential goal in image processing and computer vision [8]. Roughly speaking, edges may be defined as pixel intensity discontinuities or jumps in intensity within an image. Edges typically occur on the boundary between two different regions in an image. The goal of edge detection is to extract important features from the edges of an image. Identifying and localizing these edges are a low-level task in a variety of applications such as 3D reconstruction, shape recognition, image compression, and image enhancement and restoration. Intensity changes in images are usually caused by geometric events in the form of discontinuities in surface orientation, object depth, and/or surface color and texture. Many classical edge-detection methods rely on the image derivatives to detect points that lie on edges. For example, the Sobel method finds edges using the Sobel approximation to the derivative, while the Canny method finds edges by looking for local maxima of the image gradient.

Multiscale methods based on wavelets have been successfully applied to the analysis and detection of edges [9, 10]. Despite their success, wavelets are, however, known to have a limited capability in dealing with directional information. By contrast, shearlets are particularly designed to deal with directional and anisotropic features typically present in natural images, and have the ability to accurately and efficiently capture the geometric information of edges. As a result, the shearlet framework provides highly competitive algorithms for detecting both the location and orientation of edges, and for extracting and classifying basic edge features such as corners and junctions.

The shearlet approach has similarities with a number of other methods in applied mathematics and engineering to overcome the limitations of traditional wavelets. These methods include contourlets, complex wavelets, ridgelets, curvelets, and bandelets. In contrast to all these methods, the shearlet framework provides a unique combination of mathematical rigidness and computational efficiency when addressing edges, optimal efficiency in dealing with edges, and computational efficiency. In addition, the continuous formulation of shearlets is particularly well suited for designing an implementation for the purpose of edge analysis and detection. Furthermore, the shearlet approach is

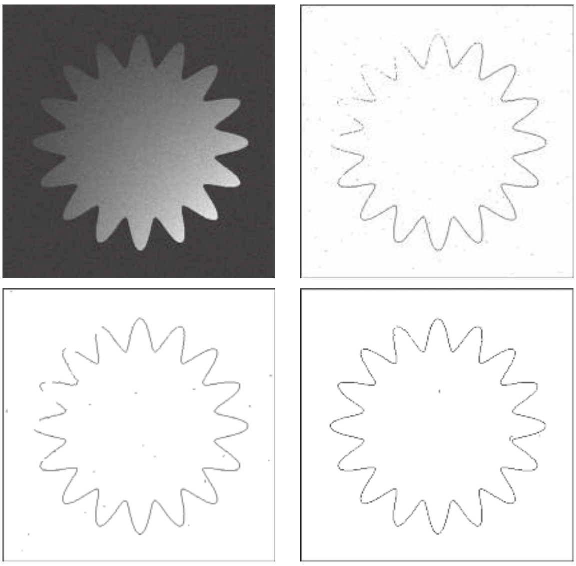

Figure 1.4 Edge detection results (clockwise from top to bottom): noisy shape; Sobel result; wavelet result; shearlet result.

Figure 1.5 Edge detection results (clockwise from top to bottom): noisy Lena image; Sobel result; wavelet result; shearlet result.

based on a simple and rigorous mathematical theory which accounts for the geometrical properties of edges [11]. Numerical tests show that the shearlet-based edge-detection method is very effective at detecting edges compared to other well-established edge detectors, as depicted in Figures 1.4 and 1.5.

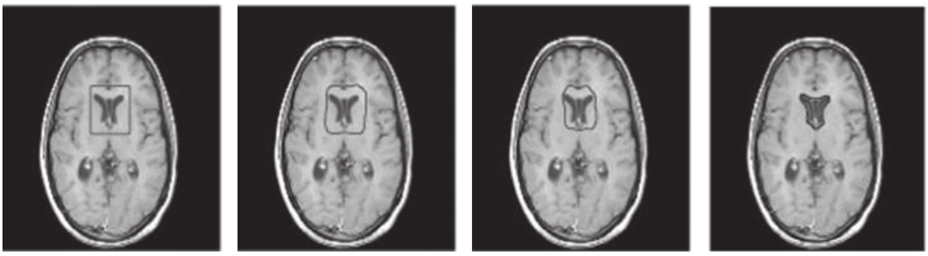

Figure 1.6 Segmentation results using curve evolution.

1.3.2 Image segmentation

Image segmentation is the process of partitioning an image into multiple parts in an effort to identify objects or other relevant information in the image [12]. Image segmentation methods may be divided into several categories, including thresholding techniques, edge-based methods, region-based techniques, curve evolution methods, and clustering approaches. In the curve evolution methods, for example, curves are evolved in an image from initial locations in order to detect object boundaries (represented by closed curves) within the image [13, 14]. These methods are typically formulated in the calculus of variations framework [15]. Figure 1.6 shows an example of medical image segmentation using curve evolution [16]. From left to right, notice that the initial contour (rectangle) flows in both directions toward the boundary.

1.3.3 Diffusion tensor imaging

Diffusion tensor imaging is a state-of-the-art magnetic resonance imaging technique for analyzing the underlying white matter structure of the brain and investigating the microstructure of biological tissue, especially in the presence of fibrous structures [17]. At each voxel of a diffusion tensor image, the water diffusion anisotropy and preferred orientation can be measured and represented by a symmetric second-order tensor. The orientation of the resulting diffusion tensor field represents the orientation of fiber bundles, and hence diffusion tensor imaging is considered an ideal choice for studying and inspecting white matter metabolism in the brain. By detecting the orientation of water molecules in white matter, diffusion tensor imaging enables the study of white matter alteration across populations and provides a helpful tool for brain growth research.

An important prerequisite for these studies is nonrigid image registration, which refers to the process of aligning two or more images of the same scene that were subject to elastic or nonrigid transformations so that their details overlap accurately [18]. Images are usually registered for the purpose of combining or comparing them, enabling the fusion of information in the images. Extending nonrigid image registration from scalar images to diffusion tensor images is, however, a challenging task, not only because of the multidimensionality of diffusion tensor images, but also due in large part to the requirement of keeping diffusion tensor orientation consistent with the anatomy after image transformation.

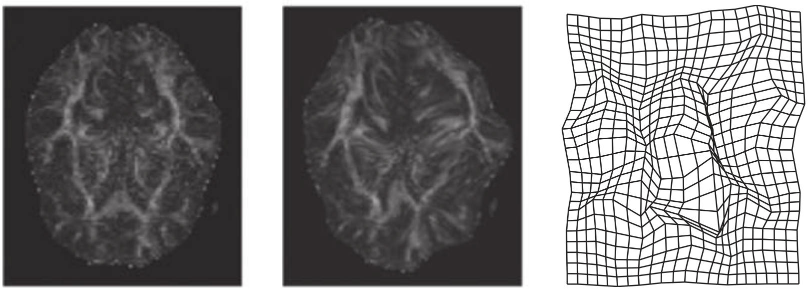

Figure 1.7 Fixed image (left); moving image (center); deformation field (right).

Roughly speaking, the image registration (also called alignment) problem may be formulated as a two-step process. The first step is to define a similarity measure that quantifies the quality of alignment between the fixed image and the transformed moving image. The moving image is a result of a deformation field applied to the fixed image, as shown in Figure 1.7. The second step is to develop an efficient optimization algorithm for maximizing the similarity measure in order to find the optimal transformation parameters of the deformation field.

1.3.4 Surface denoising

The great challenge in image processing and computer graphics is to devise computationally efficient algorithms for recovering images and surfaces (3D shapes) contaminated by noise and preserving their geometrical structure. With the increasing use of 3D scanners to create 3D models, there is a rising need for robust and efficient surface denoising techniques to remove undesirable noise from the data. Even with high-fidelity scanners, the acquired models are invariably noisy, and therefore require filtering. Mesh or surface denoising refers to the process of recovering a 3D model contaminated by noise, as shown in Figure 1.8. The challenge of the problem of interest lies in faithfully recovering the original model from the observed model, and furthering the estimation by making use of any prior knowledge/assumptions about the noise process.

Generally speaking, surface denoising methods may be classified into two major categories: isotropic and anisotropic. The former techniques filter the noisy data independently of direction, while the latter methods modify the diffusion equation to make it nonlinear or anisotropic in an effort to preserve the sharp features of the surface. Most of these nonlinear methods are inspired by anisotropic-type diffusions in the image processing literature. The most commonly used mesh denoising method is the Laplacian flow, which repeatedly and simultaneously adjusts the location of each mesh vertex to the geometric center of its neighboring vertices [19]. Although the Laplacian smoothing flow is simple and fast, it produces, however, the shrinking effect and an oversmoothing result. To circumvent these limitations, several robust surface denoising techniques have been proposed in the literature, including the mean, median, bilateral, and vertex-based

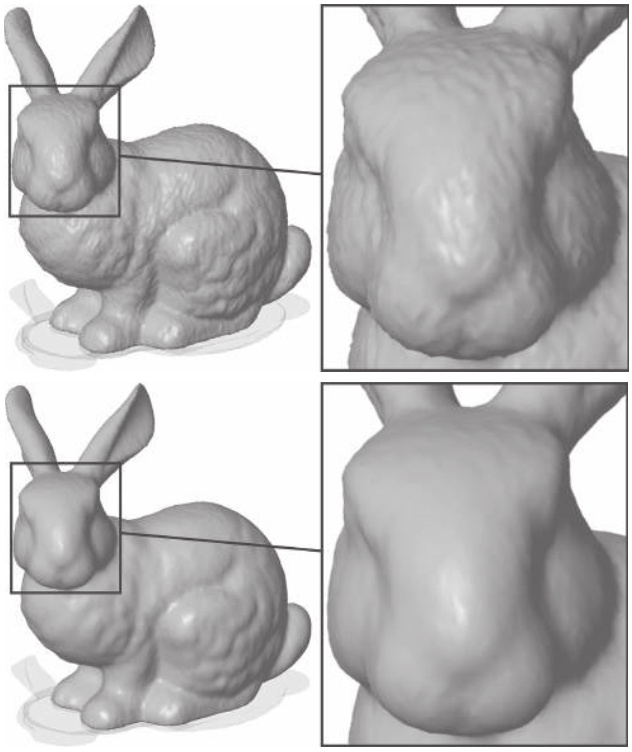

Figure 1.8 Noisy 3D bunny model (top row); denoised 3D bunny model (bottom row).

anisotropic diffusion filters. The vertex-based anisotropic diffusion uses geometric insight to help construct an efficient and fast 3D mesh smoothing strategy to fully preserve the geometric structure of the data [20]. Intuitively, the smoothing effect of the vertex-based diffusion flow may be explained as follows. In flat regions of a 3D mesh where the vertex gradient magnitudes are relatively small, the diffusion flow is reduced to the heat equation which tends to smooth more, but the smoothing effect is unnoticeable; around the sharp features of the 3D mesh where the vertex gradient magnitudes are large, the diffusion flow tends to smooth less, and hence leads to a much better preservation of the mesh geometric structures.

1.3.5 Surface compression

Compression of images and 3D shapes has long been the central theme of image processing and computer vision. Its importance is increasing rapidly in the field of computer graphics and multimedia communication because it is difficult to transmit digital information efficiently over the internet without its compression. Surfaces of 3D objects consist of geometric and topological information, and their compressed representation is an important step towards a variety of computer graphics applications, including indexing, retrieval, and matching in a database of 3D models. Figure 1.9 shows the

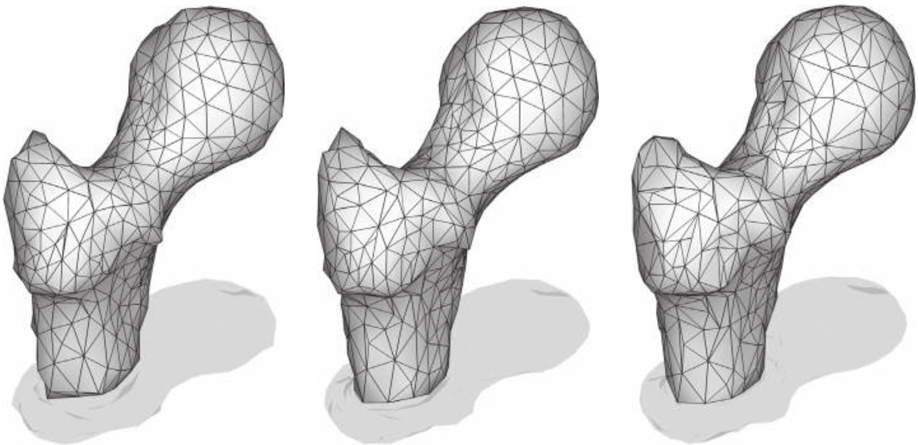

Figure 1.9 Spectral compression of a 3D model: 500 basis functions (left); 250 basis functions (center); 125 basis functions (right).

spectral compression of a 3D shape using different numbers of the umbrella operator's basis functions. The umbrella operator is a transform that basically moves each vertex of a triangle mesh to the centroid of its neighbors. The primary motivation behind using the umbrella operator is to encode a 3D shape into a parsimonious representation by retaining the smallest eigenvalues/eigenvectors of the mesh umbrella matrix. The largest eigenvalues and the associated eigenvectors are essentially discarded, thereby reducing the dimensionality of the new basis [21]. In other words, most of the energy is concentrated in the low-frequency coefficients.

1.3.6 Shape skeletonization

The first step in 3D shape matching usually involves finding a reliable skeletal graph that will encode efficiently the 3D shape information. Skeletonization aims at reducing the dimensionality of a 3D shape while preserving its topology. Skeletons have been widely used in image processing and computer vision applications [22–25].

It is widely believed that perception utilizes crucial topological characteristics of objects. Recent neuro-imaging studies, together with behavioral studies, provide strong evidence supporting the notion that topological properties (such as a genus, or connectedness) are primitives of visual representation in humans. To address the potentially complex topology of a 3D object (e.g. holes, breaks in objects) we call upon Morse theory, which essentially unfolds the topology of an object, by studying the critical points (i.e. maximum, minimum, or saddle points) of a Morse function defined on its surface [26]. These critical points, along with the regular points, fully capture the topology of a 3D object by a skeletal graph, referred to as a Reeb graph. A major advantage of a topological graph representation of an object is its ability to match objects by parts, thereby enabling a transition from a global to a localized correspondence between shapes (e.g. mechanical parts, manufactured solids). Reeb graphs have thus proven to be

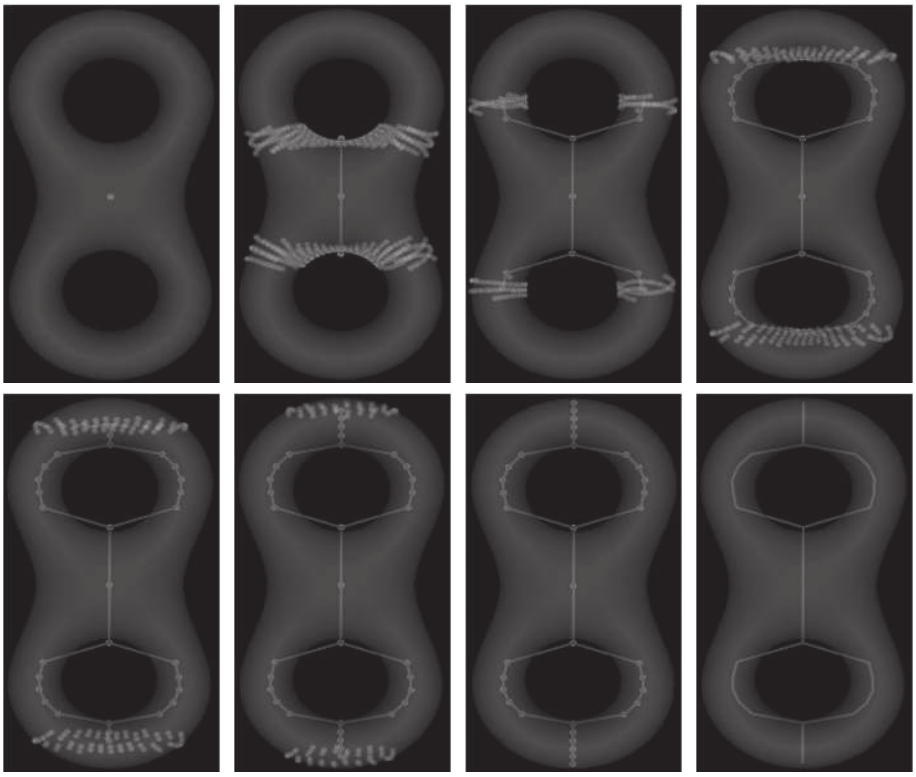

Figure 1.10 Skeletonization of a double torus.

instrumental in globally describing 3D shapes and in representing their topology [26]. Specifically, the nodes on such graphs represent critical points of a Morse function defined on a 3D shape/object surface, while the edges capture the topologically homogeneous parts, which correspond to "primitive" shapes. Figure 1.10 shows the work flow of a skeletonization process using the Morse distance function [23, 25]. It is important to understand how the distance function relates to topology and Morse theory. The intersection of a 3D object with growing spheres results in connected components on the object surface. The critical points of the surface, which define its topology, correspond to distance levels at which there is a change in the number of connected components. This change in the number of connected components reflects the changes in topology of a 3D shape, particularly, branching and merging of connected components as illustrated in Figure 1.10.

1.3.7 Shape recognition

In light of the latest software, hardware, and computing advancements, 3D technology has grown beyond being a buzzword. Today, 3D technology has become an essential part of the modern lifestyle and is gaining momentum rapidly, from consumer demand for in-home 3D television experiences to far-reaching positive implications for