Post-processing techniques for making reliable measurements from curve-skeletons

Robert S. Bradley, Philip J. Withers *

Henry Moseley X-ray Imaging Facility, School of Materials, The University of Manchester, Oxford Road, Manchester M13 9PL, UK

ABSTRACT

Interconnected 3-D networks occur widely in biology and the geometry of such branched networks can be described by curve-skeletons, allowing parameters such as path lengths, path tortuosities and cross-sectional thicknesses to be quantified. However, curve-skeletons are typically sensitive to small scale surface features which may arise from noise in the imaging data. In this paper, new post-processing techniques for curve-skeletons are presented which ensure that measurements of lengths and thicknesses are less sensitive to these small scale surface features. The techniques achieve sub-voxel accuracy and are based on a minimal sphere-network representation in which the object is modelled as a string of minimally overlapping spheres, and as such samples the object on a scale related to the local thickness. A new measure of cross-sectional dimension termed the modal radius is defined and shown to be more robust in comparison with the standard measure (the internal radius), while retaining the desirable feature of capturing the size of structures in terms of a single measure. The techniques are demonstrated by application to trabecular bone and tumour vascular network case studies where the volumetric data was obtained by high resolution computed tomography.

© 2016 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

1. Introduction

Many biological structures contain features which form a 3-D interconnected network, such as blood vessels and nerve fibres within a tissue [1,2], air pathways in the lungs [3], trabecular bone [4] and the root systems of plants [5]. The geometry of these networks can be characterised in terms of branch connectivity, path lengths and cross-sectional thicknesses, which in turn influence the structural properties as well as the transport of gasses or liquids. These characteristics can be derived from 3-D volumetric data (e.g. from x-ray computed tomography (CT) scans [6]) by representing the structure as a curvilinear skeleton which corresponds to the 'local centres' or 'backbone' of the structure [7]. However, the measurements of path length and cross-sectional thickness using this approach can be particularly sensitive to small scale surface features, which may be inherent in the object or due to noise in the volumetric data.

In this paper, new user-independent and fast post-processing techniques for curve-skeletons are described which enable more reliable estimates of path lengths, tortuosities and cross-sectional thicknesses to be obtained. The techniques are particularly suited to voxelised data obtain by tomographic scanners and are based on a minimal sphere-network representation, in which the curve-skeleton is sampled on a scale related to the local thickness of the object. As such the object is represented as a string of minimally overlapping spheres. The techniques achieve sub-voxel accuracy and are validated using simulated data sets with varying noise levels and by application to shape repository test objects. The efficacy of the techniques is then considered by application to two case studies for which the data has been obtained by x-ray computed tomography. In the first case study (trabecular bone), the robustness to different image noise levels is considered, and in the second (tumour vasculature) a comparison is made to a minimum-cost path algorithm.

2. Background and related work

An important challenge when characterizing networks is to provide a curvilinear skeleton of the structure that describes the 'backbone' of the interconnections. The medial axis transform (MAT) is commonly used for this purpose. The medial axis is defined as the locus of the centres of maximal inscribed spheres [7,8], and as such the object is represented by an infinite union of spheres, all of which touch the object surface in at least 2 points. The medial axis is derived from the distance map, which gives the shortest distance from any voxel in the object to its boundary.

* Corresponding author.
  E-mail addresses: robert.bradley-2@manchester.ac.uk (R.S. Bradley), p.j.withers@manchester.ac.uk (P.J. Withers).

Points on the medial axis are located at local maxima in the distance map corresponding to singular points (discontinuities in the gradient of the distance map). The distance map also gives a simple measure of the local thickness of the structure (e.g. pore space) at a given point on the medial axis, i.e. the largest diameter sphere that will fit within the structure at that point. The medial axis is well known to be sensitive to small scale surface features and, in particular, protrusions can lead to extra branches in the medial axis, whereas depressions can lead to a local change in the position of the medial axis. Furthermore, only in the case of tubular structures will the medial axis be composed of 1-dimensional lines; for general 3-D shapes it will be composed of 2-D sheets (manifolds) which form 'surface' skeletons [9,10].

A curvilinear skeleton (or curve-skeleton) of a 3-D object is a more recent, but less well mathematically defined, concept in comparison to MAT. Cornea *et al.* [11] described a set of desirable properties that a curve-skeleton should possess, including the following: be 1-dimensional; be centred within the object; be homotopic to the original object; be invariant under isometric transformations; be robust to noise; be smooth. There are a wide variety of algorithms to calculate curve-skeletons from meshes, point clouds, or voxelised data, which possess the desirable properties to a greater or lesser degree. An overview of some algorithms is given in [11–13].

Robustness to small scale surface features is a key requirement if reliable measurements are to be made from curve-skeletons. It is often the low spatial frequency component of the skeleton which is of interest, particularly when the surface detail is primarily due to noise, or if it varies across the structure or between structures. Robustness can be defined in relation to 3 characteristics, namely: (1) stability of the number and connectedness of branch points and branches; (2) smoothness of the segments that make up the skeleton; (3) stability of the local thickness measure associated with each point in the skeleton. Much research has been carried out to address (1) via pruning algorithms based on a measure of the saliency of each branch [12,14,15], or for simple tubular structures by constructing the shortest path through the skeleton [16,17]. Recent research has focused on both (1) and (2) with a range of algorithms in literature being shown, largely visually, to be reasonably robust to noise. However, the algorithms typically have a number of input parameters which directly or indirectly affect the robustness so that the degree of smoothing is therefore somewhat user dependent [18–27]. Similarly, those curve-skeletons derived from meshes or surface approximations (see algorithms reviewed in [11,28] for example) will be affected, to differing degrees, by the choices made when generating the mesh or surface approximation from volumetric data, which generally include the amount of surface smoothing and the number of surface elements. Furthermore, different algorithms can produce curve-skeletons with different branch points, connectivities and branch locations. Examples of this are given in [11,12,28]. Therefore the most robust algorithms may not generate the curve-skeleton with the most desirable form for a given application. Finally, algorithms differ in their complexity, such that it may be impractical to apply some robust algorithms to large data sets (e.g. the potential field approach can take 1–2 orders of magnitude longer to compute than other algorithms [11]).

Recent promising approaches to address (1) and (2) that are particularly suited to 3-D medical imaging data are based on minimum-cost path algorithms (MCPA) [3,29–32]. The approaches involve two main steps; the first involving calculation of a cost function (via the fast marching method for example) from a start voxel to all other voxels in the structure. In the second step, the geodesic path from an end point to the start point is determined by gradient descent in the cost function, with sub-voxel resolution being possible by interpolation. The cost function weighting is chosen to ensure the geodesic paths are attracted to local centre points in the object, while being less sensitive to surface detail. A key advantage is that the approach can be applied directly to the volumetric data without segmentation, provided a suitable cost function can be specified. However, the approach has a number of disadvantages, including the fact that a suitable cost function may be complex to specify, few points may actually lie on the medial axis and it may be difficult to automatically specify the start and end points of a complex network. Care must also be taken to ensure that the computed paths do not take undesired shortcuts through the object (e.g. see [3,33]). Finally, the approach can be slow for large data sets, such as those obtained by high resolution 3-D imaging (e.g. microCT data sets typically contain 8×10^9 voxels or greater), where there are many possible paths through the structure [29].

Robustness in terms of the local thickness measure (3) has received less attention. Saha *et al.* [34] used the fuzzy distance transform to measure trabecular bone thickness from 3-D voxelized data sets. The fuzzy distance transform was sampled along the skeleton of the object's support, and they demonstrated that this measure was robust down to a signal-to-noise ratio as low as ~ 5. However, the technique is reliant on a fuzzy membership of each voxel to the object under consideration, which may not be straightforward to define, depending on the contrast between the object and background. Coeurjolly developed an alternative approach to capture the distribution of thicknesses based on the power diagram [9]. The mean values of these measures (over the object) were shown to be stable to noise, and hence are useful in capturing overall (global) statistics. Rolland-Nevière *et al.* extended the shape diameter function to develop a robust thickness estimate for 3-D objects bounded by a surface mesh [35]. The algorithm returns a scale-dependent thickness measurement on the object's surface mesh.

Approaches based on MCPA do not give a direct measure of local thickness within the object. Instead, Schena and Favretto [30] calculated the cross-sectional area on planes perpendicular to the path direction, from which a simplified description in terms of equivalent diameters could be calculated. The equivalent diameter is calculated from the cross-sectional area, A, as $2(A/\pi)^{0.5}$. While the equivalent diameters capture aspects of the cross-sectional geometry in terms of a single number, they can be misleading as cross-sections with widely differing aspect ratios can give the same value. By contrast, the medial axis approach provides an unambiguous measure of thickness, being the diameter of the largest sphere which will fit into the structure at each point on the skeleton.

3. Curve-skeleton post-processing techniques based on a minimal sphere-network representation

3.1. Motivation

In computer graphics, sphere-trees are hierarchical representations of 3-D objects made by a union of overlapping spheres, with the hierarchy given by the level of detail considered [36,37]. This volume filling representation is useful for fast collision detection and computing intersection volumes. This type of representation has also been used for segmentation of curves from 2-D images, in a technique termed 'pearling' [38]. The curves are represented as overlapping discs from which a smooth centre-line can be extracted. Here we construct a 3-D sphere-network representation by sampling the curve-skeleton with minimally overlapping spheres, this being the smallest subset of the curve-skeleton which maintains the union of spheres and also spans the object. This minimal sphere-network applied to curve-skeletons