Using grey-level and distance information for medial surface representation of volume images

Stina Svensson, Ingela Nyström
Centre for Image Analysis
Lägerhyddsvägen 17, 75237 Uppsala, SWEDEN
stina@cb.uu.se, ingela@cb.uu.se

Carlo Arcelli, Gabriella Sanniti di Baja
Istituto di Cibernetica, CNR
Via Campi Flegrei 34, 80078 Pozzuoli, Napoli, ITALY
car@imagm.cib.na.cnr.it, gsdb@imagm.cib.na.cnr.it

Abstract

A medial surface representation of a grey-level volume image is computed. The foreground is reduced to a subset topologically equivalent to the initial foreground and mainly consisting of surfaces centred within regions having locally higher intensities, here, regarded as more informative. This result is obtained by combining distance information with grey-level information. A surface skeleton is first computed, where excessive shortening is prevented by a regularity condition defined on the distance transform. The structure of the surface skeleton is then simplified by removing some peripheral surfaces, so obtaining the desired medial surface representation.

1. Introduction

The number of imaging devices generating 3D (volume) images and the applications utilizing these images are increasing. A 3D image is generally grey-level and hard to transform into a bi-level image without loss of important information. This paper presents a shape representation, where the grey-level distribution of the image is in focus.

Volume images contain large amounts of data, but the information relevant for shape analysis is actually carried by a limited number of voxels. Thus, there is an interest for representation schemes with reduced dimensionality, e.g., the skeleton. Skeletonization of bi-level 2D images has been widely investigated. For grey-level 2D images, skeletonization has been faced thoroughly only during the last decade, [1, 2, 3, 7]. For bi-level 3D images, where all foreground voxels have the same grey-level and, hence, the same relevance, a number of algorithms are available producing the desired result: the skeleton preserves the topological properties of the foreground and, in some cases, allows its reconstruction. For grey-level 3D images, where the foreground is characterized by several grey-levels, the contributions are few, [6, 9]. In this case, representing the foreground in terms of entities of lower dimensionality, surfaces and curves, is not a straightforward task, as the relevant information is not evenly distributed. Here, we refer to image domains where the relevant information is gathering in correspondence with the foreground subsets with locally higher intensities, so that surfaces and curves should be found mainly therein. This is the case, e.g., for MRA (Magnetic Resonance Angiography) images, where the foreground is the blood vessels, which are characterized by high intensities.

We outline a method to reduce the foreground of a grey-level 3D image to a medial surface representation. This representation is topologically equivalent to the initial foreground and is mainly constituted by surfaces (and curves) centred within regions having locally higher intensities. Our medial surface representation is obtained via a process that we call 3D grey-skeletonization. The set resulting from it, the *grey-skeleton*, is processed to maintain the most significant surfaces (and curves), originating the desired grey-medial surface representation (*grey-MSR*). Our definition of a grey-skeleton (and, hence, of a grey-MSR) of a 3D image descends from the notions of the skeleton of a grey-level 2D image, [2], and the skeleton of a bi-level 3D image, [11].

2. Notions and notations

We consider a grey-level volume image G. Any voxel v in G has 26 neighbours in the $3 \times 3 \times 3$ set of voxels centred on v: 6 sharing a face, 12 sharing an edge, and 8 sharing a vertex with v. The voxels are denoted n_i, $i = 1, \ldots, 26$, and constitute the 26-neighbourhood of v. The set consisting of the voxels in the 26-neighbourhood of v, except for the vertex neighbours of v, is the 18-neighbourhood. The voxels in G are assumed to have one of a finite number of integer values, g_k, $k = 0, \ldots, n$, ordered increasingly. The value of v is denoted $g(v)$.

A *region* R with grey-level g_k is a maximal connected set of voxels all having grey-level g_k. A *cavity* is a region R, whose adjacent regions (i.e., the regions having at least one n_i in the cavity) have grey-levels greater than g_k.

The *background* consists of all regions with grey-level g_0 together with all cavities. For simplicity, we will assume that there exists a region with grey-level g_0 that includes the planes of the image with minimum and maximum indices, i.e., these planes are part of the background in the image. The *foreground* is the union of all regions with grey-level $g_k, k = 1, \ldots, n$. We choose 26-connectedness for the foreground and, hence, 6-connectedness for the background.

In an image having only regions with two different grey-levels, g_0 and g_1, i.e., a bi-level image, we say that a voxel v belonging to a region R with grey-level g_1 is *simple* if by changing the grey-level of v to g_0, the resulting region is topologically equivalent to R. A decision on whether v is simple or not can be taken based on the number of foreground components, i.e., the number of regions with grey-level g_1, and the number of background components, i.e., the number of regions with grey-level g_0, in the $3 \times 3 \times 3$ neighbourhood of v. In fact, v is simple if the number of foreground components in the 26-neighbourhood of v, N^{26}, is equal to 1 and the number of background components in the 18-neighbourhood of v that are face connected to v, $\overline{N}_f^{18}$, is equal to 1, [4, 10]. These two conditions will be referred to as constituting the *Simplicity Condition*.

A voxel v in a grey-level image G is simple if the above condition holds in the $3 \times 3 \times 3$ bi-level neighbourhood obtained by assigning n_i properly to either foreground or background.

The distance transform (DT) of a region is a replica of the region, where voxels are labelled with their distance from a reference set, [5]. In a grey-level image G, the reference set for a region is constituted by the regions having lesser grey-level. This is an extension to 3D of the 2D case, [2]. The distance transformation is neither effective for the region with level g_0, nor for the possibly existing cavities, since their corresponding reference set is empty. The union of the DTs of all regions is the DT of G. Here, we use a DT, where the distance between two voxels is equal to the number of steps in a shortest 6-connected path between the voxels.

When computing the DT of a region with grey-level g_k, adjacent regions with grey-levels greater than g_k are obstacles for the propagation of distance information. Thus, differently from DT computed on convex domains, the distance transformation of a region with grey-level g_k may require more than one pair of forward and backward scans of the image, or a more sophisticated recursive propagation algorithm, [8]. The DTs of all regions of G can be computed simultaneously, by keeping track of the grey-levels of the voxels in G while propagating distance information.

In the $3 \times 3 \times 3$ neighbourhood of v, voxels with distance label equal to that of v can be interpreted as border voxels separating background voxels, i.e., voxels with lesser distance labels, from internal voxels, i.e., voxels with greater distance labels. In the DT of a region, we say that a voxel v with a pair of opposite face neighbours, such that one has distance label less than v (background) and the other greater than v (internal) satisfies the *DT Regularity Condition*.

To extend the DT Regularity Condition and the Simplicity Condition to the DT of G, both distance labels and grey-levels should be taken into account. The following interpretation is done. For a voxel v, the n_i with grey-level equal to $g(v)$ and with lesser distance label, or with grey-level less than $g(v)$, are considered as background voxels. Any n_i with grey-level equal to $g(v)$ and with greater distance label, or with grey-level greater than $g(v)$, is considered as an internal voxel. The voxel v and all n_i with grey-level equal to $g(v)$ and with the same distance label as v are considered as border voxels. Border voxels and internal voxels are together interpreted as foreground voxels.

3. Method

Our idea is to first compute the grey-skeleton of the foreground, in such a way that is has the same number of components, cavities, and tunnels as the initial foreground. (Of course, the 3D grey-skeletonization is different from bi-level skeletonization, where only the distance from the border of the foreground is taken into account to obtain a centred skeleton.) Then, the grey-skeleton is further processed to obtain a medial representation that is centred within the most significant parts of the foreground, i.e., the regions with higher intensities.

Only small synthetic images are shown in this paper, since this is the only way to allow visualization at voxel resolution. In Figure 1, left, a cross section of a box-shaped foreground is shown. The box-shaped foreground, i.e., the part of the image with grey-levels between g_1 and g_n, is of size $50 \times 24 \times 24$ voxels. Grey-levels decrease from the cross section towards the back and the front (in the electronic version of this paper, grey-levels are represented by colours with increasing hue angle from red=g_1 to green=g_n).

The first phase of our 3D grey-skeletonization is to compute the DT of all regions in the image. For each region, the voxels with grey-level g_k are labelled with the distance to their closest voxel in the reference set. As remarked in Section 2, voxels placed in cavities are not reached by the

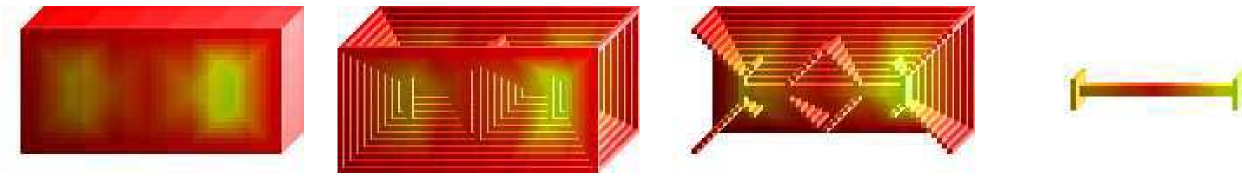

Figure 1. Box-shaped foreground: cross section; grey-skeleton and cross section; grey-MSR.

propagation of distance information and, hence, do not receive any distance label. Absence of distance label allows detection of cavities after computation of the DT of G. Voxels placed in cavities are then assigned to the background, i.e., set to grey-level g_0.

The second (and last) phase of 3D grey-skeletonization is done grey-level after grey-level. We borrow some ideas introduced in [11] for the bi-level case, to deal with grey-level images. Iterative thinning of the foreground is done, guided by the DT to identify the voxels that at each iteration constitute the border of the current foreground. In fact, voxels with different (increasing) distance labels can be interpreted as different borders found at successive thinning iterations. Thus, the distance label of a voxel can be interpreted as related to the iteration at which that voxel belongs to the border of the current foreground.

For a given grey-level g_k, voxels are examined in increasing distance label order. Each iteration is done in two steps. During the first step, the current border is identified. It includes all voxels already accepted as skeletal voxels and all voxels with grey-level g_k and distance label equal to the iteration number. Note that voxels already accepted as skeletal voxels have grey-level less than g_k, or grey-level g_k and distance label lesser than the iteration number. Voxels with distance label equal to the iteration number and with grey-level g_k are candidates for removal. To avoid creation of false cavities, we consider as candidates only voxels having a face neighbour with grey-level g_0. These voxels are all processed parallelwise and those satisfying the Simplicity Condition and the DT Regularity Condition are marked. During the second step, marked voxels are sequentially inspected and those still satisfying both conditions when visited are removed, i.e., set to g_0. Analogously to the bi-level skeletonization, marked voxels that cannot be removed are accepted as skeletal voxels (SVs). The removal process ends after all borders in all grey-levels have been processed.

What generally remains, after the second phase of 3D grey-skeletonization, is an image where the voxels either have been set to g_0 or maintain their grey-level and have been accepted as SVs. Indeed, besides these voxels, there may be voxels that still have their initial grey-level and are distance labelled, but are not accepted as SVs. This situation occurs, typically, in presence of a *cave* in the foreground. A cave is a part of a region R_k almost enclosed in

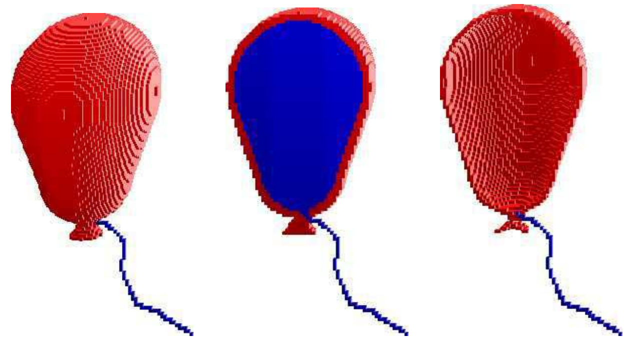

Figure 2. Grey-skeletonization in presence of caves: a balloon-shaped foreground, a cross section, and its grey-skeleton (cross section).

a region R_l, $g_l > g_k$, except for a narrow set consisting of SVs protruding from R_l. In other words, a cave is a subset of a region that is transformed into a cavity after the remaining part of the region has been assigned to the skeleton. A cave is illustrated in Figure 2, where a balloon-shaped foreground consists of two regions with grey-levels g_1 and g_2, respectively, $g_1 < g_2$ (g_1 and g_2 are shown in blue and red, respectively, in the electronic version of this paper). At the end of the removal process, the string of the balloon entirely consists of SVs, while the rest of the region with grey-level g_1 consists of voxels that are neither set to g_0, nor assigned to the skeleton. In fact, these voxels can not be removed and can not be taken as SVs, having no neighbour in the background. These voxels constitute the cave. After identification of the SVs, the voxels of the caves are set to g_0.

If *all* voxels in the current border were considered for removal, all voxels in the cave would be accepted as SVs, which is undesirable, because they do not satisfy the Simplicity Condition and the DT Regularity Condition. On the contrary, by considering for removal only the subset of the border consisting of voxels with a face neighbour with grey-level g_0, we do not need any preprocessing to detect the caves, if any. A cross section of the grey-skeleton is shown in Figure 2, right.

Voxels that are neither removed nor accepted as SVs, do not necessarily belong to caves. There might be voxels in

1051-4651/02 $17.00 (c) 2002 IEEE

pathological cases placed in thick junctions of the surface skeleton. When these voxels are ascribed to the background together with voxels located in caves, false cavities are created, at most two-voxel thick in at least one of the three main directions. To fill these false cavities, one step of an expansion/shrinking process is done. During expansion, any voxel that is neither removed, nor accepted as SV and that has at least one skeletal voxel as face neighbour, is accepted as SV. Expansion accepts as SVs all voxels in the false cavities as well as some voxels in the caves. Shrinking is then performed to remove the voxels in the caves erroneously accepted as SVs. This is done by iterative removal of SVs with a face neighbour that is neither removed, nor accepted as an SV. The expansion/shrinking process is actually done *before* the voxels in the caves are set to g_0.

The grey-skeleton, topologically equivalent to the initial foreground, is shown in Figure 1, middle, together with a cross section. Though topologically correct and with lower dimensionality with respect to the foreground, the grey-skeleton is still a too rich structure to be manageable.

At this stage, the grey-skeleton is processed to reduce it to the surfaces (and curves) regarded as the most important. To this purpose, we inspect the grey-skeleton iteratively to decide for each of its voxels whether it should be kept in the grey-MSR or not. The current edge of the grey-skeleton is identified by using the classification described in [12]. Edge voxels are inspected in increasing grey-level and, for each grey-level, in increasing distance label order, analogously to what was done for the grey-skeletonization. A voxel is removed if at least one of its n_i has equal grey-level and greater distance label, or greater grey-level, and is not necessary for topology preservation. In this way, the resulting grey-MSR will have a simpler structure with respect to the grey-skeleton, as the peripheral non 6-connected surfaces have been removed. See Figure 1, right.

Due to lack of space, we can not include other examples, but point out that for real images a preprocessing step is generally necessary before starting grey-skeletonization to reduce local grey-level variations.

4. Concluding remarks

We have described a method to compute a medial surface representation of 3D grey-level images. The method combines distance information with grey-level information. A surface skeleton is first computed, where excessive shortening is prevented by a regularity condition defined on the distance transform. This set is topologically equivalent to the initial foreground. To simplify the structure of the surface skeleton, some of its peripheral surfaces are removed to obtain the desired grey-medial surface representation. This set is mainly constituted by surfaces (and curves) centred within regions having locally higher intensities, which we regard as more significant.

If the grey-level distribution is such that no cavities exist, then the foreground is analogous to a one-level solid volume object. In this case, a curve skeleton can be computed from the grey-skeleton, e.g., by using the bi-level method based on detection of curves and junctions in a surface skeleton, [12]. We will further investigate the use of such grey-level curve skeletons.

Our work is still in progress. In this paper, we have tested the method mainly on synthetic images, because our primary interest was to verify how geometrical/topological features were maintained during the process. We will work with real images in a near future.

References

[1] K. Abe, F. Mizutani, and C. Wang. Thinning of gray-scale images with combined sequential and parallel conditions for pixel removal. *IEEE Trans. on Systems, Man, and Cybernetics*, 24:294–299, 1994.

[2] C. Arcelli and G. Ramella. Sketching a grey-tone pattern from its distance transform. *Pattern Recognition*, 29(12):2033–2045, 1996.

[3] C. Arcelli and L. Serino. Regularization of graphlike sets in gray-tone digital images. *Int. Journal of Pattern Recognition and Artificial Intelligence*, 15(4):643–657, 2001.

[4] G. Bertrand and G. Malandain. A new characterization of three-dimensional simple points. *Pattern Recognition Letters*, 15:169–175, 1994.

[5] G. Borgefors. On digital distance transforms in three dimensions. *Computer Vision and Image Understanding*, 64(3):368–376, 1996.

[6] P. Dokládal, C. Lohou, L. Perroton, and G. Bertrand. A new thinning algorithm and its application to extraction of blood vessels. In *Proc. of Biomedsim'99*, pp. 32–37, ESIEE Noisy-le-Grand, France, 1999.

[7] S.-W. Lee and Y. J. Kim. Direct extraction of topographic features for gray scale character recognition. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 17(7):724–729, 1995.

[8] J. Piper and E. Granum. Computing distance transformations in convex and non-convex domains. *Pattern Recognition*, 20(6):599–615, 1987.

[9] S. M. Pizer, D. Eberly, D. S. Fritsch, and B. S. Morse. Zoom-invariant vision of figural shape: The mathematics of cores. *Computer Vision and Image Understanding*, 69(1):55–61, 1998.

[10] P. K. Saha and B. B. Chaudhuri. Detection of 3-D simple points for topology preserving transformations with application to thinning. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 16(10):1028–1032, 1994.

[11] S. Svensson. Detecting a D^6 surface skeleton using iterative thinning. In *Proc. of Swedish Symposium on Image Analysis (SSAB'00)*, pp. 37–40, Halmstad, Sweden, 2000.

[12] S. Svensson, I. Nystr¨om, and G. Sanniti di Baja. Curve skeletonization of surface-like objects in 3D images guided by voxel classification. *Pattern Recognition Letters*, 23(12):1419–1426, 2002. Scheduled publication.