Fast and Accurate Interactive Image Segmentation in the GeoMap Framework

Hans Meine, Ullrich Köthe and Hans-Siegfried Stiehl

Cognitive Systems Group, University of Hamburg,
Vogt-Kölln-Str. 30, 22527 Hamburg, Germany
{meine,koethe,stiehl}@informatik.uni-hamburg.de

Abstract. Although many interactive segmentation methods exists,none can be considered a silver bullet for all clinical tasks. Moreover, incompatible data representations prevent multiple algorithms from being combined as desired. We propose the GeoMap as a unified representation for segmentation results and illustrate how it facilitates the design of an integrated framework for interactive medical image analysis. Results show the high flexibility and performance of the new framework.

1 Introduction

Currently, fully automatic segmentation of medical images is neither feasible nor desirable. Having a "user in the loop" is necessary from both a clinical and a legal point of view. Following the paradigm of interactive segmentation, a number of approaches were proposed which combine the cognitive abilities and medical experience of humans with the reproducable accuracy and computational power of machines. Such approaches differ in how they balance speed, ease-of-use, accuracy, reliability and other design criteria. No single method achieves the optimal balance for all classes of images or at least for all clinically relevant regions of a single image. Therefore, combinations of several methods are required.

Current toolkits (e.g. [4,11]) usually contain various segmentation algorithms, but offer only limited means to combine them on a single image. Ideally, a fully integrated tool environment would make it possible to i) switch to the most appropriate method depending on the local image content (e.g. employ edge and region detectors in the same image), and ii) reuse components of one method in another one (e.g. Canny's hysteresis thresholding within a watershed segmentation). Such combinations are currently difficult because different algorithms use incompatible data representations for almost all levels beyond the pixel matrix.

To solve this problem, a unified data representation based on sound computer science principles is needed. Therefore, we link image analysis know-how with the ideas of abstract data types and modern generic programming techniques [1]. As a first result of our research program we propose the GeoMap, a new representation based on topological maps [5] which covers the requirements of a large number of algorithms. By using GeoMaps for intermediate and final segmentation results, the combination of algorithms is made possible and, in fact, easy to realize. Our new approach can be implemented very efficiently, achieving interactive response times even on a low-cost PC.

2 Requirements for Method Integration

In order to achieve the high degree of method integration briefly discussed above, we have to i) identify common characteristics of all segmentation methods, ii) cope with their differences, and iii) develop a unified representation to work with. In general, segmentation approaches can be classified into edge-based and region-based ones. One key to unification is to exploit the duality of boundaries and regions. Duality requires the inverse of an edge representation to be a region representation and vice versa. Unfortunately, the standard implementations of popular algorithms do not possess this property. By relying on duality, one can easily switch back and forth between edge- and region-based approaches as needed and without information loss. Segmentation then amounts to finding the most significant boundaries among a large number of candidates. In region-based approaches, edge significance is derived from statistics of the adjacent regions, whereas most edge-based ones employ locally computed edge strength measures.

Below we will specify our proposed representation as an abstract data type (ADT) whose capabilities are derived from a requirement analysis of segmentation algorithms. We can only summarize requirements here, for details see [7]:

Basic Entities Low-level segmentation algorithms typically extract three types of features: regions, edges, and corners/junctions. These three correspond to the topological entities that are theoretically required for a consistent partitioning of the plane (in both the continuous and discrete domains). Therefore, our ADT supports three types of *cells*: faces, edges, and vertices.

Topology Queries The topological part of our representation encodes complete information about the neighborhoods of and adjacencies between all cells describing a particular partitioning of the image plane. The ADT provides convenient access to these relations.

Geometry Queries The geometric part of our representation associates shape information with the abstract topological cells. This can be used to either derive geometric cell properties or to access the cells' underlying pixels in order to compute intensity statistics.

Transformations We consider segmentation as the transformation of an initial partitioning of the image plane (possibly a trivial one where every pixel is a separate region) into the desired result. Thus our ADT offers a set of *Euler operations* which i) perform atomic transformations (e.g. removing one edge) and ii) can be composed in arbitrary order (e.g. to merge several regions).

Inverse Geometry Queries Interactive segmentation tools require a mapping from coordinates (e.g. obtained by a pointing device) to the associated cells. This mapping can also be used to visualize the current segmentation.

Application-Specific Information The ADT must support the association of additional information (like region statistics) with each cell.

Note that in our current implementation we have not yet realized the full potential of the theoretical framework, but restricted ourselves to transformations that start with an initial oversegmentation to be reduced (as discussed below).

3 The GeoMap Framework and its Implementation

In order to fulfill the above requirements, we define the GeoMap data type. The name reflects the fact that we augment the notion of an extended combinatorial map with the necessary geometry-related functionality [5,7]. The GeoMap itself provides access to a CellInfo object for each cell (either by enumerating all cells of a given type or via a cell label or pixel coordinate) and allows to invoke Euler operations. The CellInfo objects contain the necessary application-specific data, the coordinates of all pixels belonging to the cell, and a set of DartTraversers that can be used to query the topological relations of the cell.

A DartTraverser is similiar to a cursor in a word processor: it refers to a specific location during traversal of a GeoMap's cells. At any given traversal step, it is located on a vertex and points along an edge starting there. Thus, it uniquely defines a starting vertex, an edge, and a face to the left. It offers functions to i) reverse the orientation, ii) turn around the start vertex, iii) follow the contour of the left face, and iv) access the CellInfos of incident cells.

In contrast to other frameworks that provide topological data structures [4,11], consistency of the representation is automatically enforced by the GeoMap because Euler operations are guaranteed to perform only admissible transformations. This takes a major burden off the algorithm implementer and facilitates the execution of formal complexity analysis and correctness verification.

In principle, the GeoMap framework can be used to improve an arbitrary initial segmentation, but its realization is much simplified when we start with an oversegmentation. In this case, only Euler operations reducing the number of cells are needed. This can be seen as an example of the superpixel approach recently introduced by Malik [10]. We define superpixels by performing a watershed transform of the gradient image at fine scales. Several studies suggest this to be a good starting point for medical image segmentation [6,9].

Superpixel representations and the GeoMap complement each other in an ideal way: i) Computational cost is reduced, since the number of superpixels is much lower than the number of raw pixels. Moreover, *generic programming* techniques [1] can almost entirely eliminate the abstraction penalty usually imposed by topological approaches. ii) Being region-based, superpixel properties are more informative than pixel-based measurements. Additionally, the connection to the pixel plane is never lost during the segmentation process, and one can always access the original data to compute additional cell properties. iii) Edge detection can be interpreted as selecting significant edges among all superpixel boundaries, whereas region growing boils down to merging superpixels.

The restriction to simplifying Euler operators allows to store a sequence of reductions in a GeoMapPyramid which logs the operation applied in each step. This enables the user to backtrack to any previous segmentation state if the process arrived at a wrong result. Each level of the pyramid is a GeoMap, so any GeoMap-based algorithm can be applied at any time to compute additional pyramid levels. Our pyramid achieves a much finer granularity (in fact, the finest possible) than other irregular pyramid definitions (e.g. [2]), where a large number of reductions is always executed in parallel.

Fig. 1. Brainweb Image (181x217) ITK example. Simulated MR slice originally from http://www.bic.mni.mcgill.ca/brainweb/

Preprocessing:	2.5 sec.
Ventricles & outer brain contour (Canny-like):	2.5 sec.
Gray vs. White matter (Active Paintbrush):	1 min.
Outer Skull (Intelligent Scissors):	25 sec.

4 Application of GeoMaps in the Medical Domain

In the current prototypical system, we realized the following algorithms:

Canny-like edge detection Canny's algorithm [3] is based on gradient measurements like the watershed transform which we use instead of Canny's non-maxima suppression to get closed superpixel contours. Edgel linking is implicitly performed upon superpixel insertion into a GeoMap. Hysteresis thresholding then amounts to the removal of GeoMap edges based on gradient strength measurements.
In addition, our framework allows to use alternative edge significance measures on the basis of edge orientation or region similarity. Due to edge–region duality, region merging is essentially the same process.

Active Paintbrush Edges are removed when the mouse pointer crosses them. This serves to quickly merge several regions by "painting over the edges" [6].

Intelligent Scissors After selecting a seed point on a contour, a live-wire indicating the optimal edge from the seed to the current pointer position is shown in real-time. By coarsely marking a few additional points, a complete contour can be delineated. We implemented the tobogganing variant of this algorithm where candidate edges are restricted to superpixel boundaries [8].

A typical session in our framework may proceed as follows: after preprocessing and creation of the initial superpixel-based GeoMap, the user performs a number of automatic reductions by means of the Canny-like algorithm to remove insignificant edges. Then important contours are marked with the intelligent scissors tool which makes these edges immune against removal. In some areas with either low contrast or many edges, edge removals with the active paintbrush may be preferred for marking regions, whose contours are then protected by a single command. Finally, all unprotected edges are removed automatically. Throughout the process false reductions can be undone. Figs. 1 and 2 show results obtained this way, along with the time each step took. Note that we do not claim clinical correctness due to a lack of ground-truth. Note that the dynamic nature of our tools can only be fully appreciated in a software demonstration.

5 Conclusions

The GeoMap framework introduced in this paper combines image analysis with computer science concepts in order to define a new integrated representation

Fig. 2. Original CT slice (l), superpixels in ROI (m), segmentation (approx. 40sec) (r)

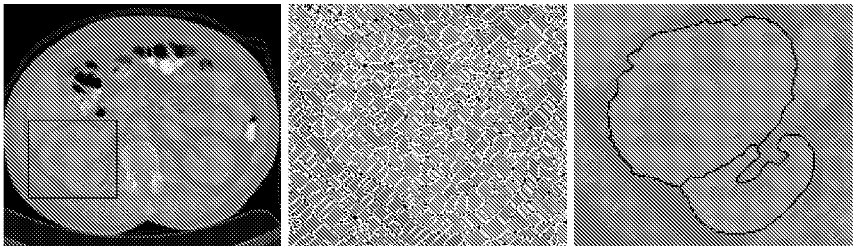

for intermediate and final results of segmentation algorithms. After algorithms have been adapted to this framework (which is not difficult and often leads to more readable code), they can be freely combined even within the same image. Thus, the most appropriate algorithm can be applied at any image location. This is especially suitable for interactive image analysis where the user guides the segmentation process and may backtrack whenever a step does not yield the desired result.

In the future many more algorithms will be integrated into our framework, especially learning based ones. It will also be necessary to collect data with clinical ground truth in order to validate our results. The extensions of our approach to 3D is also on our agenda. This promises more accurate segmentations, but requires the development of more sophisticated data structures and the application of advanced 3D interaction methods.

References

1. M.H. Austern: *Generic Programming and the STL*, Addison Wesley, 1999
2. L. Brun, W.G. Kropatsch: *Receptive fields within the Combinatorial Pyramid framework*, Graphical Models 65(1-3), pp. 23-42, 2003
3. J. Canny: *A Computational Approach to Edge Detection*, IEEE Trans. Pattern Analysis and Machine Intelligence, 8(6), pp. 679-698, 1986
4. L. Ibanez, W. Schroeder, L. Ng, J. Cates: *The ITK Software Guide*, Kitware, Inc., `http://www.itk.org/`, 2003
5. U. Köthe: *XPMaps and Topological Segmentation*, in: A. Braquelaire, J.-O. Lachaud, A. Vialard (eds.): Proc. of DGCI 2002, pp. 22-33, 2002
6. F. Maes: *Segmentation and Registration of Multimodal Medical Images*, PhD Thesis, Faculteit Geneeskunde, Katholieke Universiteit Leuven, 1998
7. H. Meine: *XPMap-Based Irregular Pyramids for Image Segmentation*, Diploma Thesis, Department of Informatics, University of Hamburg, 2003
8. E.N. Mortensen, W.A. Barrett: *Toboggan-Based Intelligent Scissors with a Four-Parameter Edge Model*, in: Proc. of CVPR'99, Vol. II, pp. 452-458, 1999
9. O.F.Olsen: *Multi-Scale Watershed Segmentation*, in: J.Sporring, M.Nielsen, L.Florack, P.Johansen(Eds.) "Gaussian Scale-Space Theory", Kluwer, 1997
10. X. Ren, J. Malik: *Learning a Classification Model for Segmentation*, in: Proc. of ICCV'03, Vol. I, pp. 10-16, 2003
11. *TargetJr Documentation*, The TargetJr Community, 1998 (project abandoned)