Medical Image Analysis 17 (2013) 1123–1136

Contents lists available at ScienceDirect

Medical Image Analysis

journal homepage: www.elsevier.com/locate/media

Delineating anatomical boundaries using the boundary fragment model

Richard V. Stebbing *, J. Alison Noble

Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, Oxford, United Kingdom

ARTICLE INFO

Article history:
Received 29 October 2012
Received in revised form 9 July 2013
Accepted 12 July 2013
Available online 23 July 2013

Keywords:
Shape model
Segmentation
Machine learning

ABSTRACT

In this paper we present a method to automatically isolate relevant anatomical boundary positions in an image using only the structure of edges. The purpose of this method is to facilitate model-based segmentation algorithms which rely on accurate initialisation and assume that the correct anatomical boundary positions are close to the current model surface.

The method is built around a weak parts-based shape model – the Boundary Fragment Model (BFM) – which represents an object by sections of its boundary. Following previous literature, we use the BFM in a boosted classifier framework to first automatically detect the object of interest. Extending previous work, we use the BFM to drive a classifier which isolates boundary candidates from spurious and irrelevant edge responses. The application of our algorithm leads to a labelled edge map which encodes the positions of (multiple) object boundaries.

By way of illustrating what is a general solution, the task of identifying the endocardium and epicardium in three-dimensional ultrasound images is completely examined, including a detailed analysis of the parameters which impact on the model construction, the structure of the learned edge response classifier, and implementation concerns. For completeness, we also demonstrate how the output boundary positions can be used in a full model-based segmentation framework.

1. Introduction

Segmentation of biomedical images is difficult because simple intensity relationships cannot be relied upon for accurate segmentation. Instead, structure and context must be used to accurately segment anatomical objects of interest. Model-based segmentation methods (Heimann and Meinzer, 2009) are a popular class of methods for two- and three-dimensional medical image analysis. They work by deforming a model of the object boundary – represented either implicitly (Osher and Fedkiw, 2001) or explicitly (Kass et al., 1988; Cootes et al., 1995) – to the anatomical boundary of interest. Anatomical boundaries are assumed to correspond to edge responses, ridges or region boundaries, all of which are determined by an *appearance model*. In the presence of noisy boundaries or absence of information, the model boundary is also deformed under a *shape model* to constrain the segmentation. The principal advantage of model-based segmentation strategies is that geometric properties such as volume and surface area are directly recoverable from the final segmentation. This is especially useful for clinical applications such as in quantitative echocardiography where estimation of left ventricle volume is used for calculation of clinical indicators like ejection fraction.

Typical model-based segmentation techniques utilise continuous optimisation algorithms to deform the model boundary which only guarantee a local optimum. Initialisation of the model boundary is therefore required to encourage the optimisation to the desired local optimum. Furthermore, *accurate* initialisation is typically required to prevent the optimisation being misled by spurious edge responses.

In this paper we present an algorithm which complements model-based segmentation strategies by automatically detecting the anatomical object of interest and isolating edge responses which correspond to its boundary. This is achieved in two steps which we refer to as *detection* and *delineation*. Both steps utilise *only* the structure of the edge responses derived from an off-the-shelf appearance model or edge detection process, captured by a weak shape model known as the Boundary Fragment Model (BFM). The BFM is a parts-based model (Felzenszwalb and Huttenlocher, 2005) where each *edge part* is a set of edge responses which describe a section of the anatomical boundary (Fig. 2). The position of each part is defined relative to the model centroid, but this position is not fixed and the flexibility is learned during the construction of the model.

The advantage of the BFM is that it is parameterised only by centroid position and scale, enabling its use for global object detection by exhaustive search. To test a given centroid position and scale, the model is "placed" in the image and the support of the edge parts is used to determine the likelihood of the test position

* Corresponding author. Tel.: +44 7826436792.
E-mail address: richard.stebbing@eng.ox.ac.uk (R.V. Stebbing).

1361-8415/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.media.2013.07.001

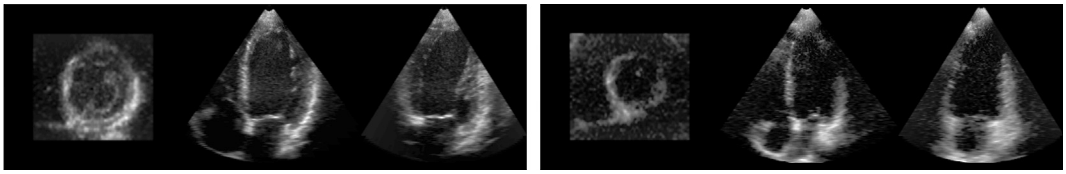

Fig. 1. Example slices from two three-dimensional echocardiograms. In each subfigure a *short-axis* slice is shown on the left, with a *four-chamber* and *two-chamber* view shown centre and right.

(Fig. 2). Following Shotton et al. (2008a) and previous work using the BFM, we use a boosted classifier to learn which parts best discriminate between different test positions.

Once detection is complete, we use the BFM to drive a separate edge response classification (delineation) step which has not been done previously. The purpose of this step is to isolate edge responses corresponding to different anatomical boundaries and suppress irrelevant spurious edge responses. To achieve this we use the BFM to derive structural features for each edge response which are then used in a Random Decision Forest (RDF) classification framework.

The work presented in this article is inspired by earlier research by Shotton et al. (2008a) and Opelt et al. (2006). The key contributions we make are generalising the BFM to three-dimensions, extending its application from object detection to object delineation, and interrogating its operation towards a common medical image analysis problem – segmenting the left ventricle in three-dimensional echocardiograms (Fig. 1).

Following this section, we provide an overview of related work (Section 2), and the details and construction of the BFM are subsequently explained in Section 3. In Section 4 object detection using the BFM is described. Next, we present the extension of the BFM to boundary delineation (Section 5). In Section 6, we present experiments to elucidate the properties of the BFM and demonstrate its operation. Finally, a model-based framework for complete segmentation is presented in Section 7.1.

2. Related work

2.1. Motivation

Model-based segmentation methods are popular in medical image analysis because they allow direct recovery of geometric properties of an object of interest. Probably the most well-known model-based segmentation strategy is Active Contours or "Snakes" (Kass et al., 1988). In this framework an object boundary is represented explicitly by a series of splines and segmentation is transformed to an energy minimisation problem over the positions of the spline control vertices. The energy function over the control vertices captures the requirement that the boundary is smooth and that it adheres to edges or region boundaries. Due to the complexity of the energy function, gradient descent or more sophisticated non-linear optimisation algorithms – which only guarantee a local energy minimum – are used to iteratively update control vertices. As a consequence, an accurate initialisation is required so that the local minimum corresponds to a solution which describes the boundary of interest.

To improve the robustness of the Active Contours framework to variation in initialisation and missing or noisy boundaries, more sophisticated shape and appearance models have been pursued which has motivated the use of alternative boundary representations. For example, Active Shape Models (Cootes et al., 1995), Active Appearance Models (Cootes et al., 2001) and its variants (Üzümcü et al., 2003; Leung and Bosch, 2007) represent the object boundary as a set of points. The points are fitted to the object boundaries on a set of training examples and some form of dimensionality reduction is used to construct a statistical shape model. Again, an iterative approach is used to fit the boundary points because of the complexity of optimising over all pose and shape parameters. At each iteration the points are moved towards the locally strongest boundary – an edge, ridge or appearance feature – which is assumed to be relevant. Next, the points are projected back into the learned shape-space.

Instead of representing the object boundary using points or splines, level-set methods (Osher and Sethian, 1988; Caselles et al., 1997; Sethian, 2003) represent the object boundary *implicitly* as the level-set of a function. The implicit representation has the advantage that additional image adherence information can be specified over easily identifiable inner and outer regions, decreasing the number of undesirable local minima and improving the quality of the final minimum solution (Lankton and Tannenbaum, 2008; Belaid et al., 2011). However, the optimisation of the energy functional using variational gradient descent leads level-set methods to suffer from the same shortcomings as explicit Active Contours.

Ultimately, regardless of boundary representation, the accuracy of a model-based segmentation strategy is bounded by the accuracy of the appearance model – edge responses detected by the appearance model must correspond to object boundaries. While many edge detection techniques in medical image analysis produce edge maps which contain the anatomically relevant edge responses, the key problems for model-based segmentation strategies have always been the presence of spurious edge responses and missing boundaries. While missing boundaries are handled exclusively by employing a shape model, we propose removing spurious edge responses by utilising knowledge about correct edge response structure. This motivates our approach of edge response classification which is complementary to improvements in domain-specific appearance models.

2.2. Object detection and voxel classification

The use of edge responses for object detection has a strong history in computer vision and a comprehensive overview is given in Shotton et al. (2008a). The focus of this section is to instead explain how the presented approach in this paper relates to current object detection and classification approaches in medical image analysis.

2.3. Discriminative object detection

Supervised classification algorithms have been utilised for many detection problems in medical image analysis. The purpose of these classification algorithms is to capture simple structure of an object of interest so that an image can be exhaustively searched. For example, AdaBoost (Friedman et al., 2000) has been used for detection of liver tumours (Pescia et al., 2008), Random Decision Forests (RDF) have been used for organ localisation in CT volumes (Criminisi et al., 2009), and support vector machines have been used for hippocampal segmentation in brain MR images (Morra et al., 2010). All of these techniques use appearance features for