Computer Vision and Image Understanding 117 (2013) 1695–1710

Contents lists available at SciVerse ScienceDirect

Computer Vision and Image Understanding

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

Elastic shapes models for improving segmentation of object boundaries in synthetic aperture sonar images

D. Bryner [a,b,*], A. Srivastava [b], Q. Huynh [a]

[a] Naval Surface Warfare Center Panama City Division, 110 Vernon Ave, Panama City, FL 32407-7001, USA
[b] Department of Statistics, Florida State University, Tallahassee, FL 32306, USA

ARTICLE INFO

Article history:
Received 9 July 2012
Accepted 1 July 2013
Available online 23 July 2013

Keywords:
Riemannian geometry
Shape analysis
Bayesian active contours
Synthetic aperture sonar

ABSTRACT

We present a variational framework for naturally incorporating prior shape knowledge in guidance of active contours for boundary extraction in images. This framework is especially suitable for images collected outside the visible spectrum, where boundary estimation is difficult due to low contrast, low resolution, and presence of noise and clutter. Accordingly, we illustrate this approach using the segmentation of various objects in synthetic aperture sonar (SAS) images of underwater terrains. We use elastic shape analysis of planar curves in which the shapes are considered as elements of a quotient space of an infinite dimensional, non-linear Riemannian manifold. Using geodesic paths under the elastic Riemannian metric, one computes sample mean and covariances of training shapes in each classes and derives statistical models for capturing class-specific shape variability. These models are then used as shape priors in a variational setting to solve for Bayesian estimation of desired contours as follows. In traditional active contour models curves are driven towards minimum of an energy composed of image and smoothing terms. We introduce an additional shape term based on shape models of relevant shape classes. The minimization of this total energy, using iterated gradient-based updates of curves, leads to an improved segmentation of object boundaries. This is demonstrated using a number of shape classes in two large SAS image datasets.

1. Introduction

An object of interest in an image can be characterized to some extent by the shape of its external boundary. It is therefore important to develop procedures for boundary extraction in problems of detection, tracking, and classification of objects in images. Certain methods for extraction make use of only the image data itself to define target boundaries while others additionally assume the availability of prior knowledge about the shape of the target to be segmented. A large body of research exists on the former approach (see for example [1,2] and papers that followed) whereas the latter is relatively less explored, with a few exceptions [3–5]. As segmentation algorithms become more sophisticated, they are tested in more difficult imaging environments of real-world scenarios where images do not have enough contrast to provide crisp, clear boundaries. One example of this scenario is when images are collected in a spectrum outside the visible domain. Here, images are typically of low contrast and contain excessive clutter, causing standard boundary extraction algorithms to fail. For instance, Fig. 8 shows some examples of synthetic aperture sonar (SAS) images that are difficult to segment automatically. Thus, it is of increasing importance that boundary extraction algorithms make use of prior knowledge about expected targets in order to help compensate for the bad data quality. Our goal is to present a method for *representing, modeling,* and *incorporating prior information* about shapes of closed curves in a boundary extraction algorithm and demonstrate its effectiveness in imaging scenarios outside the visual spectrum. We clarify that our goal is not to develop a general purpose image segmentation algorithm, but to use the contextual knowledge to determine expected object shapes and to use this information in boundary extraction. Here, one must provide prior knowledge in the form of training shapes that are representative of the target boundary desired to be segmented.

One particular example of interest to us is boundary extraction in synthetic aperture sonar (SAS) imagery in military undersea reconnaissance, although our procedure can be applied to a wide variety of applications such as medical diagnosis and infrared surveillance. The task of automatically extracting object contours in SAS imagery is challenging due to the following reasons.

* Corresponding author at: Department of Statistics, Florida State University, Tallahassee, FL 32306, USA.

E-mail addresses: dbryner@stat.fsu.edu (D. Bryner), anuj@stat.fsu.edu (A. Srivastava), quyen.huynh@navy.mil (Q. Huynh).

1077-3142/$ - see front matter Published by Elsevier Inc.
http://dx.doi.org/10.1016/j.cviu.2013.07.001

1. These images are typically low contrast due to a low signal-to-noise ratio (SNR), i.e. the background and target can be quite similar in intensity levels and boundaries are not clear. (By SNR we mean the ratio of average target pixel intensity to average background pixel intensity.)
2. The SAS imagery used here comes from a side scan sonar, where an autonomous underwater vehicle travels in a straight line and sends out a series of sound chirps from its side. Since one target side faces away from the sonar, the target is partially occluded in sound shadow, which introduces missing boundaries and causes the shape of a highlighted target to vary widely with its aspect angle.
3. The resolution of sonar images is often much lower than those obtained in the visible spectrum, resulting in relatively fewer pixels on targets.
4. Underwater imaging environments normally contain high speckle clutter due to rough seabed backgrounds.

These factors make it difficult to perform boundary extractions on SAS imagery using standard active contour and other boundary segmentation methods, and methods that can incorporate additional information are required. We will utilize manually-extracted boundaries of different known classes of sea-floor targets imaged via SAS as prior knowledge to the algorithm.

1.1. Past and present ideas in boundary extraction

The term *active contour*, or *snake*, refers to a dynamic curve that evolves to capture desired features in an image domain. There are two broad categories of active contour methods: parametric and geometric, each with their own set of advantages and drawbacks. Parametric active contours are explicitly defined parameterized curves, and the forces that drive the snake evolution are applied to the curve directly; whereas, geometric active contours are implicitly defined as zero level sets of higher-dimensional functions across the entire image domain. Here, we provide a brief survey of past and present work in these two areas.

1.1.1. Parametric active contours

Many of the past and current parametric snake models are based on the ideas presented in the seminal paper of Kass et al. [6]. They define an energy function comprised of two parts, external and internal energy, whereby an explicitly parameterized snake evolves towards local minima corresponding to desirable solutions (edges, boundaries, etc.). The external energy is made from the pixel intensity values across the image domain, and the internal energy is designed as a regularization term that applies a smoothing force to maintain continuity of the curve itself. In order to allow for a more general initialization, Cohen and Cohen in [7] introduce an internal balloon force derived from the finite element method that instead of drawing a snake inward towards a boundary edge, it inflates the snake outward. Xu and Prince in [8] introduce a new type of external force called Gradient Vector Flow (GVF), which is not based on the negative gradient of an energy function, but rather it diffuses the vectors from the Gaussian smoothed image gradient by solving a pair of decoupled linear PDE's. Li et al. [9] propose an Edge Preserving GVF (EPGVF) that maintains the benefits of GVF while improving the ability to detect weak edges. It also allows for topological change in a parametric active contour setting, a quality that arises much more naturally in geometric models.

1.1.2. Geometric active contours

A geometric active contour is described by a curve evolution equation that does not depend on the explicit parametrization of the contour but rather on the intrinsic geometric properties or quantities of the contour, such as curvature and normal vector, that are independent of parametrization. Most geometric active contour methods stem from the ideas presented in Osher and Sethian's work on front propagation with curvature-dependent speed [10]. They represent a front, or an interface, as the zero level set of a higher-dimensional signed distance function called the level-set function. The front is evolved according to hyperbolic conservation laws acting on the level set function itself, where the velocity field is defined by its curvature (geometric term) as well as possible external flows (advection term). A desirable property of this method is that it can naturally handle topological change, i.e. splitting and merging, of the evolving front with no additional effort. Malladi et al. [2] and Caselles et al. [11] independently applied Osher and Sethian's ideas to the problem of boundary extraction in images in order to make use of topological change in the segmentation of not just one but potentially many object boundaries. In addition to any curvature smoothing term, Malladi et al. define a speed function based on image data to be applied to the propagating front, which provides a halting criterion at potential boundaries. Caselles et al. use a similar approach but consider the advection term to drive active contours towards minimal distance curves or geodesics in a Riemannian space derived from the image, a popular method called "geodesic active contours." Kichenassamy et al. [12] modify slightly the work defined in [2,11] by deriving their model from basic differential geometric principles. Common to all geometric active contour models is the notion of evolution according to Euclidean curve shortening, which defines the gradient direction in which the Euclidean perimeter shrinks the fastest. This velocity vector is determined by the geometric heat equation and is equivalent to the inward normal direction scaled by the signed curvature [13,14]. Such an evolution is a very desirable component of a snake model because it simultaneously shrinks the curve and smooths it without inducing any cross-overs. Li et al. [15] also describe a level set approach for image segmentation.

1.2. Active contours with shape priors

In many applications the image information alone is seldom enough to drive the contour towards the desired target boundary, and it is necessary to incorporate prior knowledge about the type of shape. Sonar imagery is a good example of a scenario where without an energy term to help guide the snake to a set of high probability prior shapes, background clutter, occlusion, and overall inhomogeneity of the target intensity can give rise to false segmentation results. For this reason, many current methods employ a Bayesian framework to apply a *shape prior*, an energy term based on the statistics of a set of training shapes, to the active contour model. The nature of the shape prior depends on the type of active contour model in use, parametric or geometric. In a parametric model, the shape prior is a statistical model on closed or open contours in $\mathbb{R}^2$; whereas, in a geometric model, the shape prior is a statistical model on level-set functions, i.e. surfaces, in $\mathbb{R}^3$ or higher.

Past Bayesian methods have been applied almost exclusively to geometric models where most efforts follow the ideas presented in Leventon et al. [16]. Here, a set of signed distance functions are obtained from previously known training shapes, and an arithmetic mean is computed. PCA is then performed in $\mathbb{L}^2$ space to create a multivariate Gaussian distribution on the subspace defined by the top n principal components. Fang and Chan in [17] incorporate such a Gaussian density for use as a shape prior in a geodesic active contour framework based on [11]. Cremers et al. in [3] and Rousson and Cremers in [18] improve on [16] by applying a kernel density estimation technique to model general distributions of level-set functions beyond that of just Gaussian. They modify the Mumford–Shah based segmentation in [1] to include the shape