Pattern Recognition Letters 34 (2013) 833–845

Contents lists available at SciVerse ScienceDirect

Pattern Recognition Letters

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

Robust and efficient object segmentation using pseudo-elastica

Matthias Krueger [*,1], Patrice Delmas, Georgy Gimel'farb

Department of Computer Science, The University of Auckland, Auckland 1142, New Zealand

ARTICLE INFO

Article history:
Available online 3 January 2013

Keywords:
Object segmentation
Image segmentation
Second-order energy
Curvature regularity
Active contour
Elastica

ABSTRACT

A new object segmentation method based on second-order energy minimisation is proposed. It is called *pseudo-elastica* as it relates to the classic Euler's elastica but resulting contours cannot be expected to converge towards continuous elastica if the resolution is increased. Comparing to prior works, our segmentation technique can be easily applied to both closed contours and open contours with fixed endpoints, and its computational complexity, $O(N \log N)$, is significantly lower. The efficiency is increased by extending the idea of bidirectional Dijkstra-type search to second-order energies and incorporating heuristics with some sacrifice in exact energy minimisation. Our pseudo-elastica generalises the classic first-order path-based schemes to second-order energies while maintaining the same low complexity. Experiments suggest that it scores similar or better results and usually requires considerably less user input than the state-of-the-art approaches. The algorithm can be made anisotropic in order to allow corners in the contour.

1. Introduction

Segmentation of images into specific non-overlapping regions is one of the major computer vision problems, which has been intensively explored for a long time. Besides the graph cut- (e.g. Boykov et al., 2001; Boykov and Kolmogorov, 2003) and the level set-based techniques (e.g. Caselles et al., 1997; Chan and Vese, 2001), one important family of algorithms exploits weighted shortest paths. Fischler et al. (1981) arguably were the first to extract object contours of interest by computing shortest paths in graphs, whose vertices are pixels. Later, Cohen and Kimmel (1997) developed this approach further by incorporating the fast marching method (Sethian, 1996) to escape metrication errors, being intrinsic to graph-based discrete optimisation methods and leading to inaccurate segmentation for certain graphs.

Cohen and Kimmel (1997) were also the first to succeed in computing globally optimal active contours (AC). The AC concept, introduced by Kass et al. (1988), is to segment an image by aligning a curve evolving under suitable regularity constraints with a goal region boundary. This technique, as well as its main successor, the geodesic active contour (GAC) (Caselles et al., 1997), guides the curve evolution by gradient descent minimisation of particular energy functionals. The obtained contours are, generally, only locally optimal and therefore often converge to spurious edges.

The above (Cohen and Kimmel, 1997; Fischler et al., 1981) and further methods (Appleton and Talbot, 2005; Boykov and Kolmogorov, 2003) for finding globally optimal ACs share a common property: the first-order energy functionals, depending basically on the weighted arc length. Such techniques are by definition biased towards short curves and often result in unsatisfactory segmentation. Appleton and Talbot's method is a noteworthy exception with a scale-invariant energy, yet even scale-invariant first-order methods have been shown to prefer small regions (cf. Schoenemann and Cremers, 2007).

For example, an attempt to interpolate endpoints of an incomplete circle in Fig. 1(a) with a first-order shortest path yields the boundary in Fig. 1(b). A more sophisticated approach is to incorporate the curvature κ into the energy functional in order to enforce smooth transitions. Such considerations lead very naturally to the elastica energy of a sufficiently smooth curve Γ in $\mathbb{R}^2$:

$$E_{el}(\Gamma) = \int_{\Gamma} (\alpha\kappa^2 + \beta)\, ds, \tag{1}$$

where κ is the curvature of Γ, and α and β are positive constants. The elastica as minimiser of this energy were first studied in depth by Euler (1744); for a review of some of its properties see e.g. (Shah, 2002). Since the curvature term κ favours circular structures, the elastica energy of Eq. (1) conforms better to intuitive interpolation of boundary pieces missed (cf. Mumford, 1994). Accordingly, our paper proposes an efficient user-guided image segmentation algorithm based on weighted versions of the classic elastica energy. As shown in Fig. 1(c), it interpolates the incomplete circle more satisfactorily. This property of second-order energies has already led to many successful applications in edge grouping (e.g. Elder

* Corresponding author.
E-mail addresses: m.krueger74@gmail.com (M. Krueger), p.delmas@auckland.ac.nz (P. Delmas), g.gimelfarb@auckland.ac.nz (G. Gimel'farb).

[1] While doing this research M. Krueger was with the University of Auckland. He is now with the Aduno Group, Zurich, Switzerland.

0167-8655/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.patrec.2012.12.017

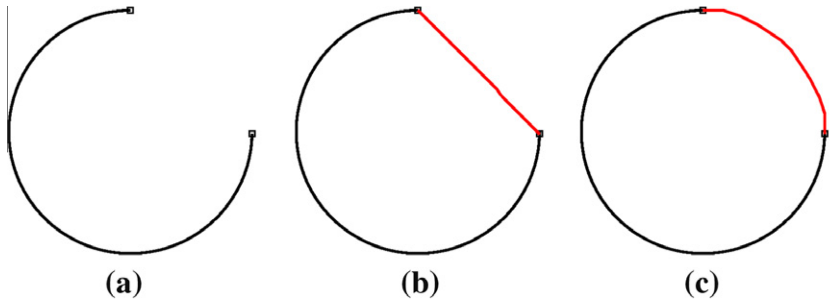

Fig. 1. Missing section of an incomplete circle (a); interpolated by a first-order shortest path (b); and an open pseudo-elastica (c) with directional constraints given at both endpoints.

and Zucker, 1996; Mahamud et al., 1999; Shashua and Ullman, 1988; Thornber and Williams, 1995; Wang et al., 2005), image impainting (e.g. Ambrosio and Masnou, 2003; Chan et al., 2002), and segmentation (Zhu and Chan, 2007; Schoenemann and Cremers, 2007; Sundaramoorthi et al., 2009; Windheuser et al., 2009; El-Zehiry and Grady, 2010a,b). We compute approximate global minimisers of such elastica-type energies fast by generalising a well-known bidirectional shortest path scheme (e.g. Nicholson, 1966; Pohl, 1971), to the second-order case. In contrast to known similar works (Schoenemann and Cremers, 2007; Windheuser et al., 2009; El-Zehiry and Grady, 2010a,b) our approach is significantly more efficient and can extract both closed contours and open contours with fixed endpoints (see Section 2 for more details and further related work). By its computational complexity of $O(N \log N)$, the proposed core algorithm can be considered as a generalisation of first-order shortest path algorithms such as (Dijkstra, 1959; Sethian, 1996) and the related image segmentation methods(Cohen and Kimmel, 1996; Fischler et al., 1981). Yet a few heuristics have to be applied in order to achieve such computational efficiency, therefore the computed contours only approximate the global optimum.

Like the algorithms in (Schoenemann and Cremers, 2007; El-Zehiry and Grady, 2010b) our technique is edge-based, i.e. it incorporates only information from the object boundary. Therefore it can readily be applied to images where region-based approaches fail, namely if the object and the background have a similar colour/intensity distribution. This is often the case for e.g. ultrasound images in medical imaging.

Note that unlike the discrete elastica proposed by Bruckstein et al. (2001) the pseudo-elastica developed in this paper cannot be expected to converge towards the continuous elastica. The algorithm clearly favours straight line segments and this does not change if the grid resolution is increased. Moreover, due to the second-order energy, the solutions found by the proposed bidirectional search scheme will in general not be optimal. As, nevertheless, in most cases the computed contours qualitatively resemble the continuous elastica (cf. Mumford, 1994; Bruckstein et al., 1996), we call them pseudo-elastica. Some experiments of our algorithm in the case of classic elastica curves are described in Section 4.4 below.

The paper is organised as follows. Section 2 reviews the known dynamic programming- and curvature-based segmentation approaches. The segmentation energies studied in the paper are introduced in Section 3, and algorithms that search for curves approximating global energy minimisers are proposed in Section 4. These core algorithms are integrated into a user-guided framework for object segmentation in Section 5, and Section 6 presents experimental results on several images, including quantitative comparisons with state-of-the-art techniques. The discussion and conclusions are given in Section 7.

2. Related Prior Work

2.1. Active Contours and Dynamic Programming

The segmentation energies used by our technique can be seen as second-order regularisations of the classic GAC functional

$$E_{GAC}(\Gamma) = \int_{\Gamma} f(s)\, ds \tag{2}$$

so that the GACs (Caselles et al., 1997) as well as the classic ACs, introduced in the seminal paper by Kass et al. (1988) certainly relate to our approach. Both these techniques use gradient descent techniques to find solutions that are in general only locally optimal.

We determine suitable contours which approximate the global minimisers, by applying discrete optimisation rather than curve evolution. Discrete approaches to image segmentation have a long history. After early approaches using dynamic programming (e.g. Montanari, 1971; Martelli, 1976) Fischler et al. (1981) were arguably the first to interpret the optimal boundary as a shortest path in a graph whose vertices correspond to the image pixels. Thereby they significantly improved the efficiency compared to the classic DP approaches. Shortly after the AC was introduced by Kass et al. (1988), Amini et al. (1990) proposed to optimise ACs with DP. Similarly to the variational approach, an initial polygon close to an object of interest has to be provided by the user. Then at each iteration the movements of the nodes minimising the AC's energy are computed with the DP.

Cohen and Kimmel (1997) adapted later and enhanced the idea by Fischler et al. (1981) by replacing the graph search with the fast marching method (Sethian, 1996) to exclude the graph-related metrication errors and find globally optimal GACs with given endpoints. Appleton and Talbot (2005) applied Cohen and Kimmel's technique in a curved product space in order to compute closed globally optimal GACs in 2D images.

Schoenemann and Cremers (2007) were the first to propose an algorithm delivering globally optimal solutions for a second-order image segmentation energy. However, the efficiency is poor because a large product graph is used, resulting in computation times of several minutes to hours. Windheuser et al. (2009) suggested a related algorithm computing the shortest path in a four-dimensional layered product graph. This high dimensional graph also results in a low computational efficiency. Later, Schoenemann et al. (2009) proposed region-based image segmentation based on linear programing, however this approach also suffers from a high running time. Recently, it was extended to multi-label problems in (Schoenemann et al., 2011), where the authors reported running times of several hours even on comparably small pictures. Strandmark and Kahl (2011) reduced the number of constraints in the optimisation problem in Schoenemann et al., 2009, thus decreasing