Pattern Recognition 48 (2015) 3293–3309

Contents lists available at ScienceDirect

Pattern Recognition

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

Image segmentation with arbitrary noise models by solving minimal surface problems

Daniel Tenbrinck [a,*], Xiaoyi Jiang [b,c,d]

[a] GREYC, UMR 6072 CNRS, École Nationale Supérieure d'Ingénieurs de Caen, France
[b] Department of Mathematics and Computer Science, University of Münster, Germany
[c] European Institute for Molecular Imaging, University of Münster, Germany
[d] Cells-in-Motion Cluster of Excellence (EXC 1003), University of Münster, Germany

ARTICLE INFO

Article history:
Received 14 February 2014
Received in revised form
12 January 2015
Accepted 14 January 2015
Available online 28 January 2015

Keywords:
Segmentation
Minimal surface problem
Thresholding
Chan–Vese model
Level set methods
Convex optimization
Physical noise

ABSTRACT

Segmentation is one of the fundamental tasks in computer vision applications. For natural images there exist a vast amount of sophisticated segmentation methods in the literature. However, these standard methods tend to fail in the presence of non-Gaussian image noise, e.g., in biomedical imaging or astronomy. In this paper we propose an adequate variational segmentation model for segmentation of images perturbed by arbitrary noise models without adding a-priori assumptions about the unknown physical noise model. For this, we discuss the prominent minimal surface problem and two different numerical minimization schemes to solve it. The first approach efficiently computes the set of all possible minimal surface solutions for the given data via convex optimization and reduces the segmentation problem to the estimation of a proper threshold. The second approach is based on level set methods and is especially suitable for the separation of inhomogeneous image regions. The advantage of this approach is both its simpleness and robustness: the noise in the data does not have to be modeled explicitly since the image intensities are separated using histogram-based thresholding techniques. The proposed model can be interpreted as a generalization of many traditional segmentation methods which implicitly assume a perturbation by additive Gaussian noise. The superiority of this approach over standard methods such as the popular Chan–Vese formulation is demonstrated on synthetic images as well as real application data from biomedical imaging.

1. Introduction

Automatic image analysis has a long history and is becoming increasingly mature. Meanwhile, the extensive application in many scientific fields poses new challenges, due to a huge variety of imaging sensors. Different sensors, e.g., synthetic aperture radar, positron emission tomography, and medical ultrasound imaging, have different physical principles and characteristics, which often require modality-specific treatment. In particular, many important imaging sensors deliver data contaminated by signal-dependent noise, i.e., the noise variance depends on the underlying signal intensity. For example, the signal in radiograph images is determined by photon counting statistics and often described as particle-limited, emphasizing the quantized and non-Gaussian nature of the signal [1,2]. Another example is the so-called speckle noise characteristic for diagnostic ultrasound imaging, which belongs to the class of multiplicative noise [3].

On the other hand, many existing image analysis methods (implicitly) assume signal-independent additive Gaussian noise and hence their application leads to suboptimal and unsatisfying results. The recent trend in the literature is to incorporate additional knowledge about the image formation process and the physical noise model into variational optimization problems. In most cases one models the image analysis task as a statistical inverse problem and uses techniques from Bayesian inference to deduce more accurate models. Several denoising methods have been proposed to take the particular characteristics of noise models into account, e.g., see [1,3]. Attention is recently increasingly given to motion analysis on images perturbed by multiplicative noise models [4,5] and also in image segmentation the positive effect of appropriate noise modeling on the results of image segmentation is reported, e.g., see [6–8]. Especially for data with poor statistics (low signal-to-noise ratio), as often present in biomedical imaging, it is important to consider the impact of non-Gaussian noise models in the process of segmentation. In [9] image segmentation into homogeneous regions is performed by minimizing a Minimum Description Length (MDL) criterion. The probability density functions of the different regions are assumed

* Corresponding author at: Department of Mathematics and Computer Science, University of Münster, Einsteinstrasse 62, 48149 Münster, Germany.
Tel.: +49 251 8333776; fax: +49 251 8332729.
E-mail address: daniel.tenbrinck@uni-muenster.de (D. Tenbrinck).

http://dx.doi.org/10.1016/j.patcog.2015.01.006
0031-3203/

to belong to the same family. In the variational segmentation model from [10] a data-fidelity term deduced from the Poisson distribution is used instead of the usual L^2 norm as a measure of fidelity. A general region-based variational segmentation framework is proposed in [11] and realized for additive Gaussian noise, Poisson noise, and multiplicative speckle noise. In particular, it is shown that this model is a direct generalization of popular segmentation models such as the Mumford–Shah model [12] or the Chan–Vese model [13] for the special case of additive Gaussian noise. The latter segmentation method has been extended in various ways. In [14] a bias field is introduced to model image inhomogeneity (the unity bias field exactly corresponds to the original Chan–Vese model) and to deal with them during segmentation. A similar energy term is proposed in [15] to cope with segmentation of noisy images.

Although all these methods report improved image segmentation performance, they all share two fundamental disadvantages. First, since a majority of these approaches are based on variational models with non-standard data fidelity terms, the minimization of the respective (not necessarily convex) energy functionals is very cost intensive in terms of computational effort. Thus, their application for huge data or even whole series of images is rather ineligible. Second, all these methods explicitly assume a certain physical noise model for which they are designed and hence their use is limited to data which fulfills these assumptions.

1.1. Contributions

Our goal is to propose a robust and flexible, but yet simple segmentation method which is able to be used universally on a wide variety of images from different imaging fields, e.g., biomedical imaging, astronomy, or natural images. First, we discuss the well-known minimal surface problem for image segmentation and give a link to the celebrated Rudin–Osher–Fatemi (ROF) total variation denoising problem. A single minimization of the latter, which can be performed efficiently with recent algorithms, gives us a denoised approximation of the given data and even more, each level set of this approximation represents a unique global minimizer of an associated minimal surface problem. This global convex optimization avoids the problem of local minima which are frequently inherent in many classical segmentation problems.

Second, we can reduce the segmentation problem to a simple threshold problem, i.e., we have to determine the best minimal surface for image segmentation among the level sets of the approximated data. Because thresholding can be performed in real-time this enables the user to choose the best solution interactively, which is especially appealing for physicians and biologists, as it gives more control to the user. Contrary, for segmenting a whole series of images it would be beneficial to perform this thresholding step automatically to reduce the needed time effort of the specialists. As there exist many thresholding algorithms in the literature, we investigate three representatives of different thresholding concepts for their ability to adapt in the presence of arbitrary noise and hence for a large class of applications.

Finally, for cases where one wants to segment inhomogeneous image regions we propose an alternative minimization scheme based on level set methods. This approach gives more control about the segmentation topology and is especially useful in situations where a global segmentation into different intensity classes is not appropriate. In this setting we introduce a linear variant of the Chan–Vese method which is more robust under noise perturbations.

1.2. Structure of this work

We start by formulating the segmentation task mathematically in Section 2 and subsequently give two special cases of segmentation models, i.e., the minimal surface problem and the level set formulation of the popular Chan–Vese method. To motivate our work we analyze the fundamental weaknesses of the latter model in Section 3. Especially the presence of non-Gaussian noise leads to misclassification of pixels during image segmentation. To fortify the significance of this problem, we discuss in Section 4 four popular physical noise models from the literature and their impact on the data. In particular, we discuss additive Gaussian noise, Loupas (speckle) noise, Poisson noise, and Rayleigh noise. In order to deal with arbitrary noise perturbations, we subsequently propose a segmentation method based on the minimal surface problem and discuss its advantages. In Section 5 we discuss two possible minimization schemes and give implementation details. The first method utilizes the close relationship to the ROF denoising problem and we use a recent primal-dual minimization approach for the total variation based energy functional. For the case of inhomogeneous target structures we give an alternative approach based on level set methods. We evaluate different automatic thresholding methods during experiments on both synthetic and real data from biomedical imaging in Section 6. Finally, we conclude this work by a short discussion in Section 7.

This paper is an extension to our preliminary work presented in [16], which uses level set methods for segmentation of ultrasound images only. In this work we embed our approach into a more global mathematical setting and demonstrate how to obtain unique global minimizers efficiently by a new numerical minimization scheme. We demonstrate its universal applicability in various settings, i.e., for different physical noise models, on both synthetic and real biomedical imaging data. Finally, different histogram-based thresholding methods are evaluated under the impact of non-Gaussian noise.

2. Mathematical formulation of the segmentation task

Since image segmentation is a fundamental task in computer vision, there exists a vast amount of approaches in the literature. Even for single biomedical imaging application the quantity of proposed segmentation methods is huge, e.g., see the comprehensive review articles in [17–20]. However, the majority of modern segmentation approaches assemble pixels to higher-order units and incorporate spatial information about the geometry of these regions, in contrast to pixel-wise segmentation approaches discussed in the comparison study in [21]. Hence, the recent trend is towards models such as active contours, which are based on PDEs and variational formulations, e.g., see [8,11–13,22–25].

Due to the fact that region-based segmentation models are less susceptible under the impact of noise perturbations compared to edge-based methods, we concentrate on these models in the following and mathematically formulate the segmentation problem in this context. Most approaches in the literature utilize image features based on the signal intensity, e.g., the partition of the image domain into subregions with homogeneous grayscale values. Typically, these methods aim to minimize the signal variance in the respective subregions, such as the Mumford–Shah model [12] and its popular variant, the Chan–Vese segmentation method [13]. We discuss the latter approach in more detail in Section 2.2. For methods based on other features, e.g., motion or texture, we refer to [22].

In segmentation the general goal is to obtain a partition $\mathcal{P}_m(\Omega)$ of a bounded, open image domain $\Omega \subset \mathbb{R}^n$ (typically $n \in \{2,3\}$) into pairwise disjoint regions Ω_i, $i = 1, \ldots, m$, according to the properties of the given image $f : \Omega \to \mathbb{R}$. This can be expressed mathematically as follows:

$$\mathcal{P}_m(\Omega) \in \left\{ (\Omega_1, \ldots, \Omega_m) : \overline{\Omega} = \bigcup_{i=1}^{m} \overline{\Omega}_i,\ \Omega_i \cap \Omega_j = \emptyset \text{ for all } i \neq j \right\}, \tag{1}$$

where $\overline{\Omega}$ denotes the closure of Ω. Note that multiphase segmentation is in general quite challenging with respect to its numerical