Int J Comput Vis (2013) 104:241–269
DOI 10.1007/s11263-013-0621-4

Discrete and Continuous Models for Partitioning Problems

Jan Lellmann · Björn Lellmann · Florian Widmann · Christoph Schnörr

Received: 11 December 2011 / Accepted: 15 March 2013 / Published online: 11 April 2013

Abstract Recently, variational relaxation techniques for approximating solutions of partitioning problems on continuous image domains have received considerable attention, since they introduce significantly less artifacts than established graph cut-based techniques. This work is concerned with the sources of such artifacts. We discuss the importance of differentiating between artifacts caused by *discretization* and those caused by *relaxation* and provide supporting numerical examples. Moreover, we consider in depth the consequences of a recent theoretical result concerning the optimality of solutions obtained using a particular relaxation method. Since the employed regularizer is quite tight, the considered relaxation generally involves a large computational cost. We propose a method to significantly reduce these costs in a fully automatic way for a large class of metrics including tree metrics, thus generalizing a method recently proposed by Strekalovskiy and Cremers (IEEE conference on computer vision and pattern recognition, pp. 1905–1911, 2011).

J. Lellmann (✉)
Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge, UK
e-mail: j.lellmann@damtp.cam.ac.uk

B. Lellmann · F. Widmann
Department of Computing, Imperial College London, London, UK
e-mail: b.lellmann10@imperial.ac.uk

F. Widmann
e-mail: f.widmann@imperial.ac.uk

C. Schnörr
Image and Pattern Analysis Group & HCI, Department of Mathematics and Computer Science, University of Heidelberg, Heidelberg, Germany
e-mail: schnoerr@math.uni-heidelberg.de

Keywords Multi-class labeling · Segmentation · Partitioning problem · Graph cut · Convex relaxation · Variational methods

1 Introduction and Overview

The issue of how to formulate, solve, and approximate labeling problems has a long-standing history as one of the classical problems in image processing, and occurs in many applications, such as segmentation, multi-view reconstruction, stitching, and inpainting (Paragios et al. 2006). In the past decade, and in particular with the introduction of maximum flow-based methods into image analysis, a very popular approach has been to first cast the problem into the form of a Markov random field with a finite number of states, i.e., to minimize an energy function depending on a finite number of labels associated with a finite number of points in the image domain.

By doing so, one obtains a finite-dimensional problem defined over a discrete set of possible configurations. Therefore the resulting problem can be treated and analyzed as a combinatorial problem, for which a large machinery of solvers is available, most notably graph cut-based algorithms such as α-expansion.

As an alternative, in the past few years it has become more and more common to postpone the discretization until the very end, and instead to work in the function space setting as long as possible.

In particular, we will consider the following formulation, which can be seen as a continuous analogue of the finite-dimensional pairwise MRF energy. A predecessor can be found in Lie et al. (2006), while the model was proposed independently in Zach et al. (2008), Lellmann et al. (2009), and Pock et al. (2009).