Int J Comput Vis (2012) 98:1–14
DOI 10.1007/s11263-011-0491-6

Multi-label Moves for MRFs with Truncated Convex Priors

Olga Veksler

Received: 24 December 2009 / Accepted: 10 August 2011 / Published online: 7 September 2011

Abstract Optimization with graph cuts became very popular in recent years. While exact optimization is possible in a few cases, many useful energy functions are NP hard to optimize. One approach to approximate optimization is the so-called *move making* algorithms. At each iteration, a move-making algorithm makes a proposal (move) for a pixel p to either keep its old label or switch to a new label. Two move-making algorithms based on graph cuts are in wide use, namely the swap and expansion. Both of these moves are binary in nature, that is they give each pixel a choice of only two labels. An evaluation of optimization techniques shows that the expansion and swap algorithms perform very well for energies where the underlying MRF has the Potts prior. However for more general priors, the swap and expansion algorithms do not perform as well. The main contribution of this paper is to develop *multi-label* moves. A multi-label move, unlike expansion and swap, gives each pixel has a choice of more than two labels to switch to. In particular, we develop several multi-label moves for truncated convex priors. We evaluate our moves on image restoration, inpainting, and stereo correspondence. We get better results than expansion and swap algorithms, both in terms of the energy value and accuracy.

Keywords Discrete optimization · Markov random fields (MRF) · Graph cuts · Truncated convex priors

O. Veksler (✉)
Computer Science Department, University of Western Ontario, London, Canada
e-mail: olga@csd.uwo.ca

1 Introduction

Energy optimization with graph cuts (Boykov et al. 1998, 2001; Ishikawa and Geiger 1998; Kolmogorov and Zabih 2002b) is very popular in computer vision and graphics. It has been used for such diverse applications as image restoration (Boykov et al. 2001), stereo and multi-view reconstruction (Boykov et al. 1998, 2001; Kolmogorov and Zabih 2001, 2002a), motion (Wills et al. 2003), texture synthesis (Kwatra et al. 2003), segmentation (Boykov and Jolly 2001; Blake et al. 2004; Rother et al. 2006; Kolmogorov et al. 2006), digital photomontage (Agarwala et al. 2004), digital tapestry (Rother et al. 2005), image generation (Nguyen et al. 2008), computational photography (Levin et al. 2007), image completion (Hays and Efros 2007), digital panoramas (Agarwala et al. 2006).

Optimization with graph cuts is successful, largely, because either an exact minimum or an approximate minimum with certain quality guarantees is found, unlike the older optimization techniques, such as Simulated Annealing (Geman and Geman 1984) or ICM (Besag 1986). It is important, therefore, to continue seeking better optimization methods for computer vision problems. This paper is about improved optimization methods for a subclass of energy functions that is useful for computer vision.

Typically an energy minimization problem is formulated in a pixel labeling framework as follows. We have a set of sites $\mathcal{P}$, and a set of labels $\mathcal{L}$. $\mathcal{P}$ is often the set of all image pixels, and $\mathcal{L}$ is a finite set that represents the property that needs to be estimated at each site, such as intensity, color, etc. The task is to assign a label f_p to each site p so that some energy function $E(f)$ is minimized. Here f is the collection of all pixel-label assignments.

If $\mathcal{L}$ has size two and $E(f)$ is submodular (Kolmogorov and Zabih 2002b), then $E(f)$ can be optimized exactly by