Int J Comput Vis
DOI 10.1007/s11263-013-0653-9

PMBP: PatchMatch Belief Propagation for Correspondence Field Estimation

Frederic Besse · Carsten Rother · Andrew Fitzgibbon · Jan Kautz

Received: 31 January 2013 / Accepted: 5 August 2013

Abstract PatchMatch (PM) is a simple, yet very powerful and successful method for optimizing continuous labelling problems. The algorithm has two main ingredients: the update of the solution space by sampling and the use of the spatial neighbourhood to propagate samples. We show how these ingredients are related to steps in a specific form of belief propagation (BP) in the continuous space, called max-product particle BP (MP-PBP). However, MP-PBP has thus far been too slow to allow complex state spaces. In the case where all nodes share a common state space and the smoothness prior favours equal values, we show that unifying the two approaches yields a new algorithm, PMBP, which is more accurate than PM and orders of magnitude faster than MP-PBP. To illustrate the benefits of our PMBP method we have built a new stereo matching algorithm with unary terms which are borrowed from the recent PM Stereo work and novel realistic pairwise terms that provide smoothness. We have experimentally verified that our method is an improvement over state-of-the-art techniques at sub-pixel accuracy level.

Keywords Correspondence fields · Belief propagation · PatchMatch

1 Introduction

This paper draws a new connection between two existing algorithms for estimation of correspondence fields between images: belief propagation (BP; Pearl 1988; Yedidia et al. 2005) and PatchMatch (PM; Barnes et al. 2009, 2010). Correspondence fields arise in problems such as dense stereo reconstruction, optical flow estimation, and a variety of computational photography applications such as recoloring, deblurring, high dynamic range imaging, and inpainting. By analysing the connection between the methods, we obtain a new algorithm which has performance superior to both its antecedents, and in the case of stereo matching, represents the current state-of-the-art on the Middlebury benchmark at sub-pixel accuracy. The first contribution of our work is a detailed description of PM and BP in terms that allow the connection between the two to be clearly described. This analysis is largely self-contained, and comprises the first major section of the paper. Our second contribution is in the use of this analysis to define a new algorithm: PMBP which, despite its relative simplicity, is more accurate than PM and orders of magnitude faster than max-product particle BP (MP-PBP).

BP is a venerable approach to the analysis of correspondence problems. The correspondence field is parametrized by a vector grid $\{\mathbf{u}_s\}_{s=1}^n$, where s indexes *nodes*, typically corresponding to image pixels, and $\mathbf{u}_s \in \mathbb{R}^d$ parametrizes the correspondence vector at node s. We shall consider a special case of BP, viewed as an energy minimization algorithm where the energy combines *unary* and *pairwise* terms

$$E(\mathbf{u}_1, \ldots, \mathbf{u}_n) = \sum_{s=1}^{n} \psi_s(\mathbf{u}_s) + \sum_{s=1}^{n} \left[\sum_{t \in N(s)} \psi_{st}(\mathbf{u}_s, \mathbf{u}_t) \right], \quad (1)$$

with $N(s)$ being the set of *pairwise neighbours* of node s. The unary energy $\psi_s(\mathbf{u}_s)$, also called the *data term*, computes the local evidence for the correspondence $\mathbf{u}_s$. For example, if $\mathbf{u}_s = (u_s, v_s)$ is a parametrization of a two-dimensional (2D) flow field between images I_1 and I_2, then one might define a weighted patch data term (where (x_s, y_s) are the image coordinates of pixel s)

F. Besse (✉) · C. Rother · A. Fitzgibbon · J. Kautz
University College London, London, UK
e-mail: f.besse@cs.ucl.ac.uk

Published online: 20 August 2013