PMBP: PatchMatch Belief Propagation for Correspondence Field Estimation

Frederic Besse[1]
f.besse@cs.ucl.ac.uk

Carsten Rother and Andrew Fitzgibbon[2]
{carrot,awf}@microsoft.com

Jan Kautz[1]
j.kautz@cs.ucl.ac.uk

[1] University College London
London, UK

[2] Microsoft Research Cambridge
Cambridge, UK

This paper draws a new connection between two existing algorithms for estimation of correspondence fields between images: Belief Propagation [4] and PatchMatch [1]. 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 PatchMatch and belief propagation in terms that allow the connection between the two to be clearly described. Our second contribution is in the use of this analysis to define a new algorithm: PatchMatch Belief Propagation (PMBP) which, despite its relative simplicity, is more accurate than PatchMatch and orders of magnitude faster than PBP.

Belief propagation (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

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

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 $\mathbf{u}_s$. On a continuous space, a natural representation using particles presents itself, closely related to Max Product Particle BP (PBP)[3]. With each node s, we associate a set of K *particles* $P_s \subset \mathbb{R}^d$, where each particle $p \in P_s$ is a candidate solution for the minimizing correspondence parameters $\mathbf{u}_s^*$. BP is a *message-passing* algorithm, where messages are defined as functions from nodes to their neighbours. Before defining the messages, which are themselves defined recursively, it is useful to define the *log disbelief* at node s as

$$B_s(\mathbf{u}_s) := \psi_s(\mathbf{u}_s) + \sum_{t\in N(s)} M_{t\to s}(\mathbf{u}_s), \quad (2)$$

in terms of which the messages, using particles, are defined as

$$M_{t\to s}(\mathbf{u}_s) := \min_{\mathbf{u}_t\in P_t}\ \psi_{st}(\mathbf{u}_s,\mathbf{u}_t) + B_t(\mathbf{u}_t) - M_{s\to t}(\mathbf{u}_t). \quad (3)$$

We note that this definition is in terms of a continuous $\mathbf{u}_s$, not restricted to the current particle set P_s, but the minimization over $\mathbf{u}_t$ is a discrete minimization over the particles P_t. At convergence, $\hat{\mathbf{u}}_s := \operatorname{argmin}_{\mathbf{u}} B_s(\mathbf{u})$ is the estimate of the minimizer.

The **PatchMatch** algorithm (PM) [1] was initially introduced as a computationally efficient way to compute a *nearest neighbour field* (NNF) between two images. In terms of energy minimization, the NNF is the global minimizer of an energy comprising unary terms only ($\psi_{st} = 0$). PM computes good minima while being very efficient. With such a powerful optimizer, more complex unary terms can be defined, yielding another class of state-of-the-art correspondence finders, exemplified by the recent introduction of PatchMatch Stereo [2]. Using the same particle notations as BP, the set P_s are initialized uniformly at random. One PM iteration then comprises a linear sweep through all nodes in an order defined by a *schedule function* $\phi(s)$, so that s is visited before s' if $\phi(s) < \phi(s')$. At node s, two update steps are performed: *propagation* and *resampling*:

- In the **propagation** step, the particle set is updated to contain the best K particles from the union of the current set and the set C_s of already-visited neighbour candidates, where "best" is defined as minimizing the unary cost $\psi_s(\cdot)$.
- The local **resampling** step (called "random search" in [1]) perturbs the particles locally according to a proposal distribution which we model as a Gaussian $\mathcal{N}(0,\sigma)$. The second step of the PM iteration updates P_S with any improved estimates from the local resampling set, for m resampling steps.

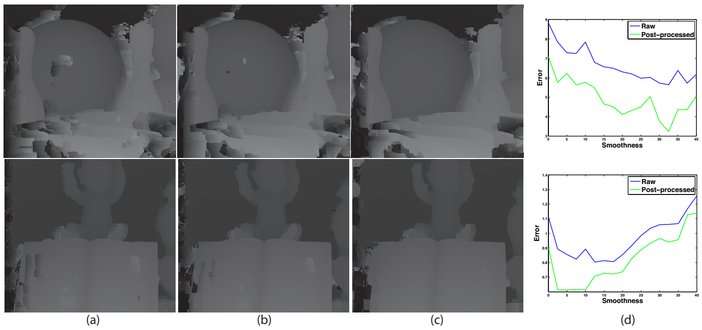

Figure 1: Evolution of the disparity map (before post-processing) with different weightings of the smoothness: (a) $\beta = 0$ (PatchMatch stereo). (b) $\beta = 5$. (c) $\beta = 17.5$. (d) Corresponding disparity error, for both raw and post-processed outputs.

After several alternating sweeps, the best particle in each set typically represents a good optimum of the unary-only energy.

PatchMatch Belief Propagation (PMBP) can be defined as a combination of the PM and PBP algorithms. We shall consider PBP our base, as the goal is to minimize a more realistic energy than PM, that is to say, an energy with pairwise terms encouraging piecewise smoothness.

First, PM resamples P_s from the neighbours of node s, while PBP's resampling is only via MCMC from the elements of P_s. The samples are evaluated using B_s, so this is a resampling of the particle set under the current belief, as proposed in PBP, but with a quite different source of particle proposals. Thus PMBP augments PBP with samples from the neighbours.

Second, PBP uses an MCMC framework where particles are replaced in P_s with probability given by the Metropolis acceptance ratio, while PM accepts only particles with higher belief than those already in P_s. This non-Metropolis replacement strategy further accelerates convergence, so it is included in PMBP.

Making these two modifications yields a powerful new optimization algorithm for energies with pairwise smoothness terms. In the case of a zero pairwise term $\psi_{st} = 0$, PMBP exactly yields PM. Conversely, running PMBP with a nonzero pairwise term is a strict generalization of PM, allowing the incorporation of an explicit smoothness control which directly addresses the deficiencies of PM while retaining its speed.

We apply our algorithm to the stereo matching case. The effect of adding a realistic pairwise term to the PatchMatch stereo algorithm under our PMBP framework can be seen in figure 1.

[1] C. Barnes, E. Shechtman, D. B. Goldman, and A. Finkelstein. The generalized PatchMatch correspondence algorithm. In *Proc. ECCV*, 2010.

[2] M. Bleyer, C. Rhemann, and C. Rother. PatchMatch Stereo—Stereo matching with slanted support windows. In *Proc. BMVC*, 2011.

[3] R. Kothapa, J. Pachecho, and E. B. Sudderth. Max-product particle belief propagation. Master's thesis, Brown University, 2011.

[4] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1988.