Maximum Persistency via Iterative Relaxed Inference with Graphical Models

Alexander Shekhovtsov[1], Paul Swoboda[2], Bogdan Savchynskyy[2,3]
[1]TU Graz, Austria. [2]Heidelberg University, Germany. [3]TU Dresden, Germany.

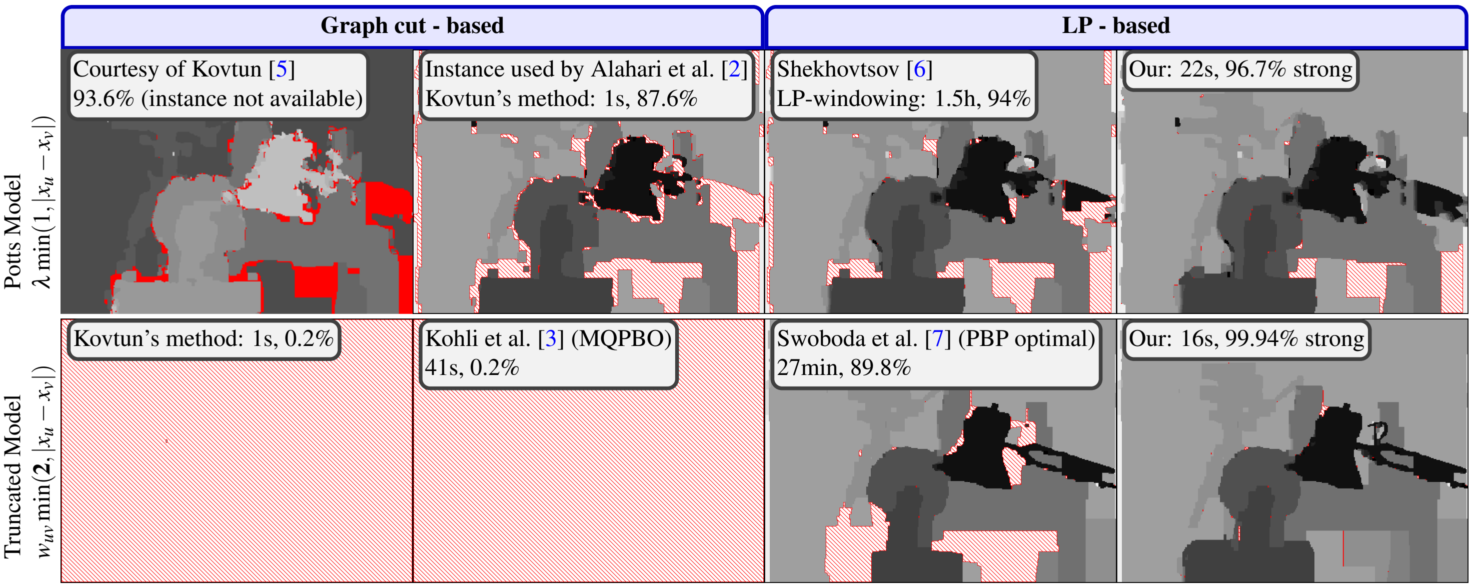

Figure 1: Progress of partial optimality methods. Top line corresponds to a stereo model with Potts interactions and large aggregating windows for unary costs used in [2, 5] (instance published by [2]). Bottom line is a more refined stereo model with truncated linear terms [8] (instance in [1]). Hashed area indicates that the optimal persistent label in the pixel is not found (but some non-optimal labels might have been eliminated). Solution completeness is given by the percent of persistent labels. Graph cut based methods are fast but only efficient for strong unary terms. LP-based methods are able to determine a larger persistent assignments but are extremely slow, prior to this work. Note, our method is set up to determine strong persistency, a partial assignment that holds for all optimal solutions, while other methods here find a part of any optimal solution.

We consider the NP-hard problem of MAP-inference for graphical models. We propose a polynomial time practically efficient algorithm for finding a part of its optimal solution. Specifically, our algorithm marks each label in each node of the considered graphical model either as (i) *optimal*, meaning that it belongs to all optimal solutions of the inference problem; (ii) *non-optimal* if it provably does not belong to any solution; or (iii) *undefined*, which means our algorithm can not make a decision regarding the label. The labels that we proved optimal or non-optimal are called *persistent*.

Key ideas:

- We build on the Maximum Persistency [6] framework, which proved that most of the existing methods for partial optimality can be explained by a simple local domination condition if only one supplies the right reparametrization of the energy function.
- Finding the maximum subset of persistent labels can be formulated [6] as a big linear program that optimizes over reparametrizations and a subset of labels deemed persistent at the same time. It is a challenging problem and large scale instances can only be addressed by a windowing technique [6] – a semi-local condition.
- We solve the same maximum persistency problem instead by iteratively solving standard LP relaxation for a series of auxiliary energy problems, similarly to the approach in [7]. We thus unite [6] and [7].

Key features of our approach:

- Invariant to reparametrization and order of labels.
- Fast approximate dual solvers can be employed without compromising correctness and global persistency guarantees.
- Requires an approximate solution to LP relaxation as a starting point.
- Can be viewed as making an approximate solver for LP-relaxation to be able to prove optimality of a part of its solution.

More specifically, we demonstrated our approach using TRW-S [4] for solving auxiliary subproblems.

Properties when subproblems are solved with TRW-S:

- Closely approximates maximum persistency LP (evaluated on small random problems).
- Fast message passing transfers to auxiliary problems.
- The method is correct using a finite number of TRW-S iterations.
- Subproblems can be solved incrementally, reusing the messages.

[1] OpenGM benchmark. http://hci.iwr.uni-heidelberg.de/opengm2/?l0=benchmark.

[2] Karteek Alahari, Pushmeet Kohli, and Philip H. S. Torr. Reduce, reuse & recycle: Efficiently solving multi-label MRFs. In *CVPR*, 2008.

[3] P. Kohli, A. Shekhovtsov, C. Rother, V. Kolmogorov, and P. Torr. On partial optimality in multi-label MRFs. In *ICML*, 2008.

[4] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *PAMI*, 28(10), October 2006. doi: 10.1109/TPAMI.2006.200. URL http://dx.doi.org/10.1109/TPAMI.2006.200.

[5] I. Kovtun. Partial optimal labeling search for a NP-hard subclass of (max, +) problems. In *DAGM-Symposium*, pages 402–409, 2003.

[6] A. Shekhovtsov. Maximum persistency in energy minimization. In *CVPR*, 2014.

[7] P. Swoboda, A. Shekhovtsov, J. H. Kappes, C. Schnörr, and B. Savchynskyy. Partial optimality by pruning for MAP-inference with general graphical models. *ArXiv e-prints*, Oct 2014.

[8] Richard Szeliski, Ramin Zabih, Daniel Scharstein, Olga Veksler, Vladimir Kolmogorov, Aseem Agarwala, Marshall Tappen, and Carsten Rother. A comparative study of energy minimization methods for Markov random fields with smoothness-based priors. *PAMI*, 30(6): 1068–1080, 2008. ISSN 0162-8828.

This is an extended abstract. The full paper is available at the Computer Vision Foundation webpage.