This CVPR2015 extended abstract is the Open Access version, provided by the Computer Vision Foundation.

Iteratively Reweighted Graph Cut for Multi-label MRFs with Non-convex Priors

Thalaiyasingam Ajanthan, Richard Hartley, Mathieu Salzmann, and Hongdong Li
Australian National University & NICTA.

While widely acknowledged as highly effective in computer vision, multi-label MRFs with non-convex priors are difficult to optimize. Even though graph-cut-based algorithms [1] have proven successful for specific problems (*e.g.*, *metric* priors), there does not seem to be a single algorithm that performs well with different non-convex priors such as the truncated quadratic, the Cauchy function and the corrupted Gaussian.

Here, we propose to fill this gap and introduce a graph-cut-based algorithm that iteratively approximates the original MRF energy with an appropriately weighted surrogate energy that is easier to minimize. We show that, under suitable conditions on the non-convex priors, our algorithm guarantees that the original energy decreases at each iteration. In particular, we consider the scenario where the global minimizer of the weighted surrogate energy can be obtained by a multi-label graph cut algorithm [3], and show that our algorithm then lets us handle of large variety of non-convex priors.

In fact, our method is inspired by the IRLS algorithm which is well-known for continuous optimization. To the best of our knowledge, this is the first time that such a technique is transposed to the MRF optimization scenario.

Let us consider a multi-label MRF with pairwise node interactions, where $\mathcal{V}$ is the set of nodes, $\mathcal{N}$ is the set of edges and each node p takes a label $x_p \in \mathcal{L}$, with $\mathcal{L}$ the set of labels. The energy of such an MRF can be expressed as

$$E(\mathbf{x}) = \sum_{p\in\mathcal{V}} \theta^u_p(x_p) + \sum_{(p,q)\in\mathcal{N}} \theta^b_{pq}(x_p,x_q) \ , \quad (1)$$

where θ^u and θ^b denote the unary potentials and pairwise potentials respectively. To apply our algorithm, we require the pairwise potentials to take the form

$$\theta^b_{pq}(x_p,x_q) = h_b \circ f_{pq}(x_p,x_q) \ , \quad (2)$$

where h_b is concave and f_{pq} is arbitrary. Then, we define the surrogate energy as

$$\tilde{E}(\mathbf{x}) = \sum_{p\in\mathcal{V}} \theta^u_p(x_p) + \sum_{(p,q)\in\mathcal{N}} w^t_{pq} f_{pq}(x_p,x_q) \ , \quad (3)$$

where $w^t_{pq} = h^s_b\big(f_{pq}(x^t_p,x^t_q)\big)$ is the supergradient of h_b. As shown in the paper, the original energy (1) can be minimized by iteratively minimizing the surrogate energy (3). In fact, the concavity of h_b guarantees that the true energy decreases at each iteration if the surrogate energy decreases.

To minimize the surrogate energy (3), we make use of the multi-label graph cut [3], which requires an ordered label set and imposes that

$$f_{pq}(x_p,x_q) = g(|x_p - x_q|) \ , \quad (4)$$

where g is a convex function. In addition, to apply max-flow, the edge weights of the multi-label graph need to be non-negative. This translates into a requirement for h_b to be non-decreasing, which comes at virtually no cost in the context of smoothness potentials in an MRF. In summary, at each iteration of our algorithm (called IRGC) we minimize

$$\tilde{E}(\mathbf{x}) = \sum_{p\in\mathcal{V}} \theta^u_p(x_p) + \sum_{(p,q)\in\mathcal{N}} w^t_{pq} g(|x_p - x_q|) \ , \quad (5)$$

where g is a convex function, and $w^t_{pq} = h^s_b\big(f_{pq}(x^t_p,x^t_q)\big)$, with h_b a concave, non-decreasing function.

In the context of computer vision problems with ordered label sets, *e.g.*, stereo and inpainting, it is often important to make use of robust estimators as pairwise potentials to better account for discontinuities, or outliers. Many such robust estimators belong to the family of functions with a single inflection point in $\mathbb{R}^+$, *e.g.*, the truncated linear, the truncated quadratic and the Cauchy function [2]. For such functions $\theta(z) = h_b \circ g(z)$, we provide a method to choose g and h_b to make the multi-label graph as sparse as possible to limit the required memory. See Fig. 1(a) for an example. Note, however, that our algorithm is not limited to the family of functions described above.

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

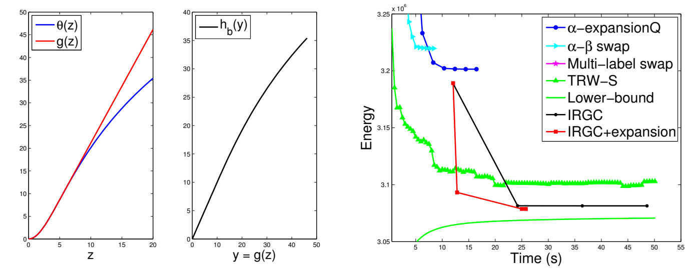

(a) θ - Cauchy function (b) Venus - stereo

Figure 1: ***(a)*** *Plots of θ, g and h_b with $\theta(z) = h_b \circ g(z)$, for θ the Cauchy function. Here g is convex and h_b is concave.* ***(b)*** *Energy vs time (seconds) for several algorithms on the Venus stereo problem with truncated quadratic prior. IRGC and IRGC+expansion outperform the other algorithms.*

	TRW-S	IRGC	IRGC+expansion
Average Quality	0.4577%	3.1116%	**0.2431%**

Table 1: *IRGC+expansion clearly yields better quality energies than TRW-S on average.*

While IRGC guarantees that the energy value decreases at each iteration, it remains prone to getting trapped in local minima. We therefore introduce a hybrid optimization strategy that combines IRGC with a different minimization technique, *e.g.*, α-expansion, and our experiments confirm that this variety in optimization is effective to overcome local minima. We refer to this algorithm as IRGC+expansion.

We evaluate our algorithms on stereo and inpainting problems and compare our results with those of α-expansion, α-β swap [1], multi-label swap [6] and TRW-S [4]. Our hybrid version consistently outperforms (or performs virtually as well as) state-of-the-art MRF energy minimization techniques. See Fig. 1(b) for an example. To evaluate the quality of the minimum energies, we followed the strategy of [5], which evaluates the relative gap between the lower bound found by TRW-S and the minimum energy of an algorithm. In Table 1, we compare these quality measures averaged over several stereo and inpainting problems.

In conclusion, with our IRGC algorithm, multi-label graph cut can be used to minimize multi-label MRF energies with arbitrary data terms and non-convex priors. Finally, IRGC really is a special case of an iteratively reweighted approach to MRF, and even continuous, energy minimization. We therefore plan to study this approach further in the future.

[1] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE T-PAMI*, 23(11):1222–1239, 2001.

[2] R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.

[3] H. Ishikawa. Exact optimization for markov random fields with convex priors. *IEEE T-PAMI*, 25(10):1333–1336, 2003.

[4] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *IEEE T-PAMI*, 28(10):1568–1583, 2006.

[5] R. Szeliski, R. Zabih, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. Tappen, and C. Rother. A comparative study of energy minimization methods for markov random fields with smoothness-based priors. *IEEE T-PAMI*, 30(6):1068–1080, 2008.

[6] O. Veksler. Multi-label moves for mrfs with truncated convex priors. *IJCV*, 98(1):1–14, 2012.