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

Thalaiyasingam Ajanthan Richard Hartley Mathieu Salzmann Hongdong Li

Australian National University NICTA

Introduction

Problem: Minimize a multi-label MRF with pairwise interactions

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

where θ^b_{pq} is non-convex.

- Non-convex priors are highly effective in computer vision.

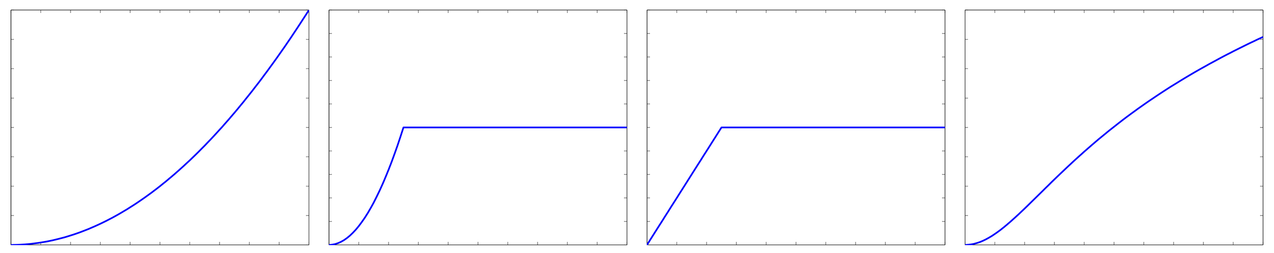

Quadratic *Tr. quadratic* *Tr. linear* *Cauchy*

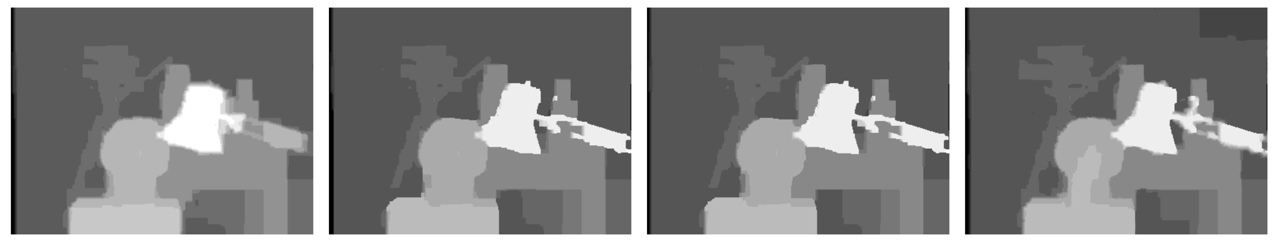

Error = 8.89% *Error = 6.49%* *Error = 6.31%* *Error = 5.04%*

Applicable graph-cut-based algorithms:

- α-expansion [2] with QPBO [1]
- $\alpha - \beta$ swap [2]
- Multi-label swap [5] (only for truncated convex priors)

Drawback: No single graph-cut-based algorithm performs well with different non-convex priors.

Contribution: Inspired by continuous optimization techniques, we introduce an iteratively reweighted graph-cut-based algorithm to minimize MRF energies.

Iteratively Reweighted Minimization

- Minimize the original energy $E(\mathbf{x})$ by iteratively minimizing a surrogate energy $\tilde{E}(\mathbf{x})$.

Algorithm

Require: $E(\mathbf{x}) \leftarrow \sum_{i=1}^k h_i \circ f_i(\mathbf{x})$ ▷ For concave functions h_i

Initialize $\mathbf{x}$

repeat

$w^t_i \leftarrow h^s_i(f_i(\mathbf{x}^t))$ ▷ Supergradient of h_i at $f_i(\mathbf{x}^t)$

$\mathbf{x}^{t+1} \leftarrow \arg\min_{\mathbf{x}} \tilde{E}(\mathbf{x}) \leftarrow \sum_{i=1}^k w^t_i f_i(\mathbf{x})$

until convergence of $E(\mathbf{x})$

return $\mathbf{x}^{t+1}$

- Either exact or approximate algorithms can be used to minimize $\tilde{E}(\mathbf{x})$.
- Guaranteed to decrease the original energy at each iteration.

Special case: Iteratively Reweighted Least Squares.

IRGC

- Minimize the surrogate energy $\tilde{E}(\mathbf{x})$ optimally using multi-label graph cut [3].

Require: For a convex function g and a non-decreasing concave function h_b

$$\theta^b_{pq}\left(x_p, x_q\right) = h_b \circ g\left(\left|x_p - x_q\right|\right)\ .$$

Algorithm:

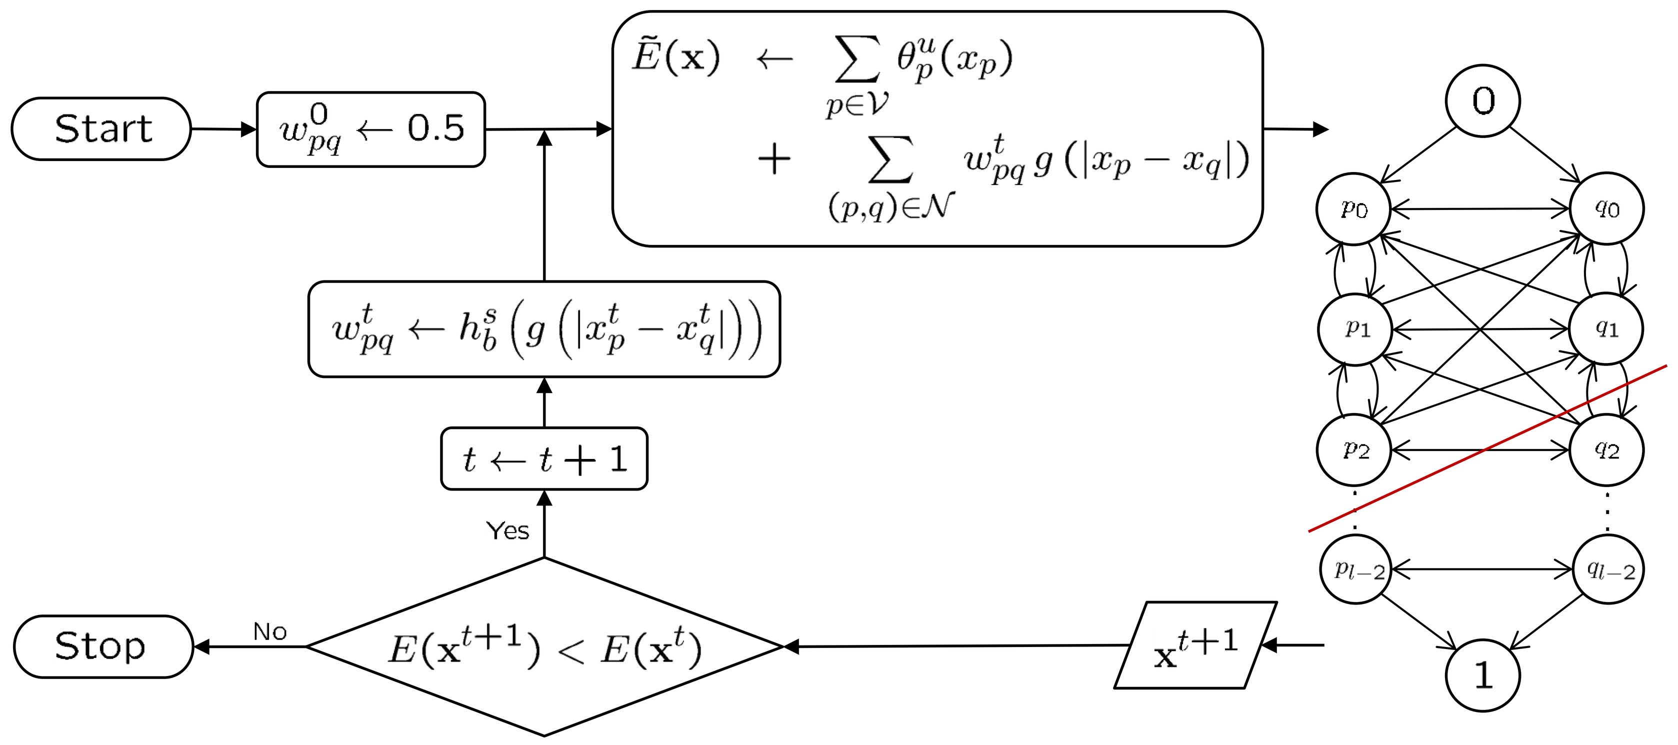

IRGC+expansion

- A hybrid optimization strategy that combines IRGC with α-expansion.

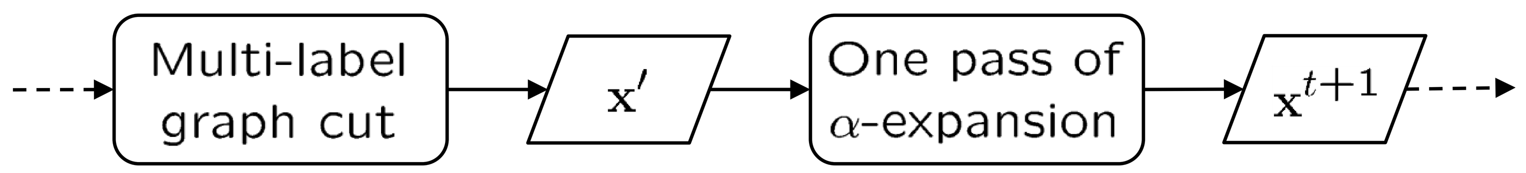

- Effective to overcome local minima.

Choice of functions g and h_b

- Consider functions $\theta(z) = h_b \circ g(z)$ with a single inflection point ($z = \lambda$) in $\mathbb{R}^+$.
- Choose g such that $g(z) = \theta(z)$ for $z \leq \lambda$ and linear otherwise.
- For such a function g, the multi-label graph requires the least amount of memory; *i.e.*, cross edge weights of the graph become zero for as many values of z as possible.

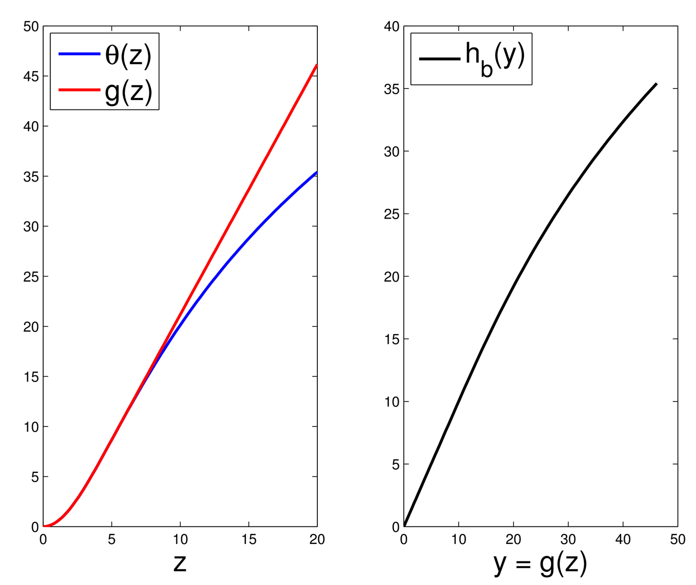

θ - Cauchy function

Experiments

- Stereo instances = 6, inpainting instances = 2.

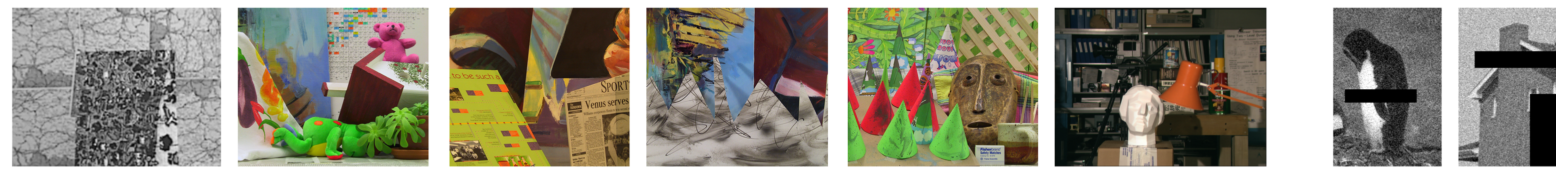

Truncated linear *Truncated quadratic* *Cauchy function* *Truncated quadratic*

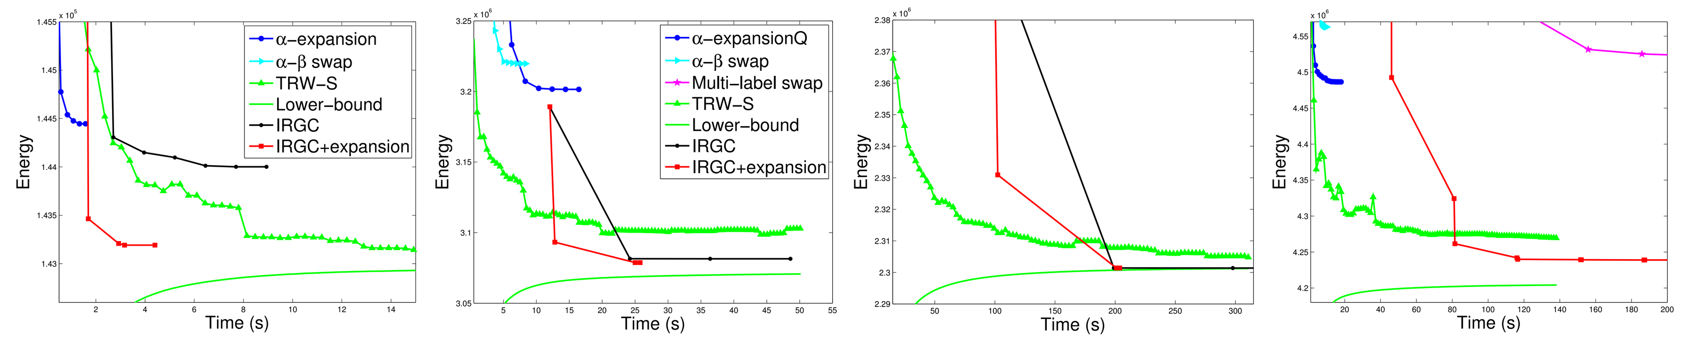

Map *Venus* *Cones* *Penguin*

Problem	α-exp. (QPBO)	α − β swap	TRW-S [4]	IRGC	IRGC+ exp.
Map	1.05%	4.59%	**0.05%**	0.74%	0.17%
Teddy	0.75%	2.40%	0.30%	1.63%	**0.21%**
Venus	4.26%	4.85%	0.91%	0.35%	**0.26%**
Sawtooth	3.42%	4.58%	0.65%	0.96%	**0.26%**
Cones	7.80%	95.11%	0.16%	**0.01%**	**0.01%**
Tsukuba	1.99%	3.45%	**0.09%**	0.47%	0.17%
Penguin	6.71%	8.53%	1.56%	11.72%	**0.83%**
House	4.59%	3.71%	0.02%	**0.01%**	**0.01%**
Average	3.82%	15.90%	0.47%	1.99%	**0.24%**

Quality of the minimum energies.

Conclusion

- Our iteratively reweighted technique provides an effective approach to minimize multi-label MRF energies with non-convex priors.
- IRGC+expansion consistently outperforms or performs virtually as well as state-of-the-art MRF energy minimization techniques.

References

[1] E. Boros and P. L. Hammer. Pseudo-boolean optimization. *Discrete applied mathematics*, 123(1):155–225, 2002.

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

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