Global MAP-Optimality by Shrinking the Combinatorial Search Area with Convex Relaxation

Heidelberg Collaboratory for Image Processing

Bogdan Savchynskyy[1] Jörg Hendrik Kappes[2] Paul Swoboda[2] Christoph Schnörr[1,2]

[1] Heidelberg Collaboratory for Image Processing, [2] Image & Pattern Analysis Group, University of Heidelberg, Germany

e-mail: bogdan.savchynskyy@iwr.uni-heidelberg.de

Abstract

- We consider energy minimization for undirected graphical models, known as MAP- or MLE-inference.
- We propose a novel method of combining combinatorial and convex programming techniques to obtain *an optimal integer solution* of the initial combinatorial problem.
- Our method enables to confine the application of the combinatorial solver to a small fraction of the initial graphical model, where the convex programming solver fails.
- The method shows superior results on a computer vision benchmark. In particular we report solving so far unsolved large scale benchmark problems and outperform in speed a state-of-the-art specialized method on Potts models.

Problem Formulation

Given the graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, associated variables $x_v \in \mathcal{X}_v$, $v \in \mathcal{V}$, and potentials $\theta_{w,x_w} \in \mathbb{R}$, $w \in \mathcal{V} \cup \mathcal{E}$, we consider the energy minimization problem

$$\min_{x \in \mathcal{X}} E(\theta, x) = \min_{x \in \mathcal{X}} \left\{ \sum_{v \in \mathcal{V}} \theta_{v,x_v} + \sum_{uv \in \mathcal{E}} \theta_{uv,x_{uv}} \right\} = \min_{x \in \mathcal{X}} \langle \theta, \delta(x) \rangle = \min_{\mu \in \mathrm{conv}(\delta(\mathcal{X}))} \langle \theta, \mu \rangle \, .$$

LP Relaxation

$$\min_{\mu \in \Lambda} \langle \theta, \mu \rangle \; : \; \Lambda = \underbrace{\left\{ \mu \geq 0 \colon \sum_{x_v} \mu_v = 1, \sum_{x_u} \mu_{uv,x_{uv}} = \mu_{v,x_v}, \sum_{x_v} \mu_{uv,x_{uv}} = \mu_{u,x_u} \right\}}_{\text{local polytope (LP)}} \supset \mathrm{conv}(\delta(\mathcal{X})),$$

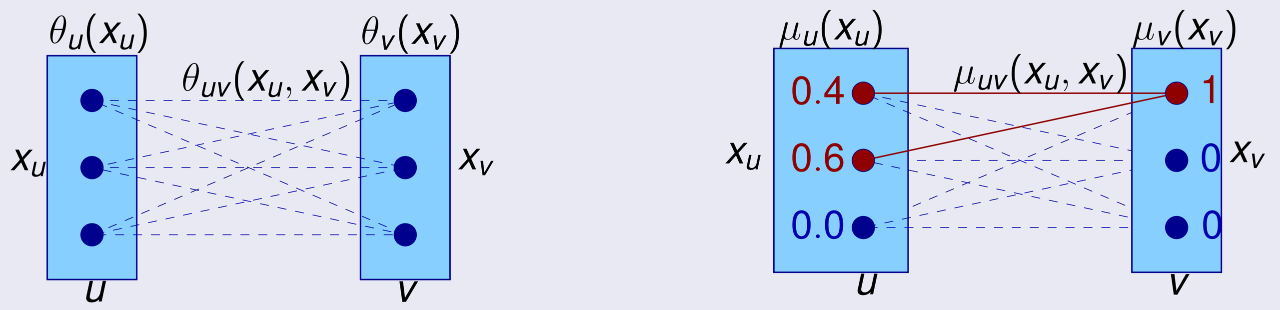

Typical Relaxed Solution

Input

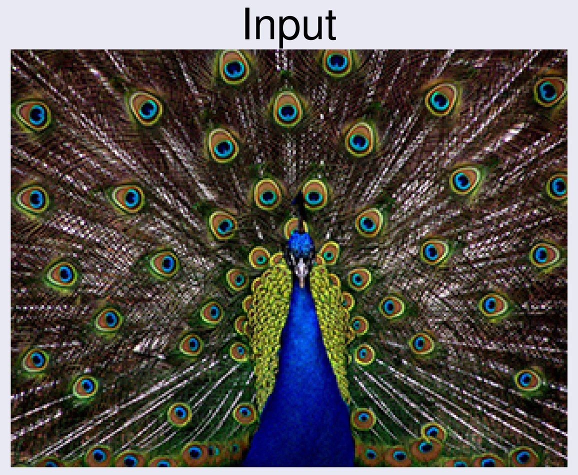

color segmentation problem

LP solution

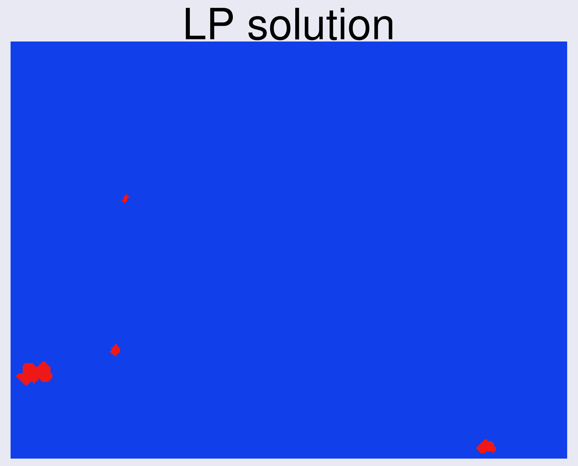

integer and fractional labels

- Is the integer part of the solution correct?
- In general - NO! In practice - mostly YES.
- How can it be exploited to find **an optimal** integer solution?

Algorithm Idea

- the strictly arc consistent nodes
- nodes with fractional labels
- border nodes
- nodes with inconsistent LP (+) and ILP (+) solutions.

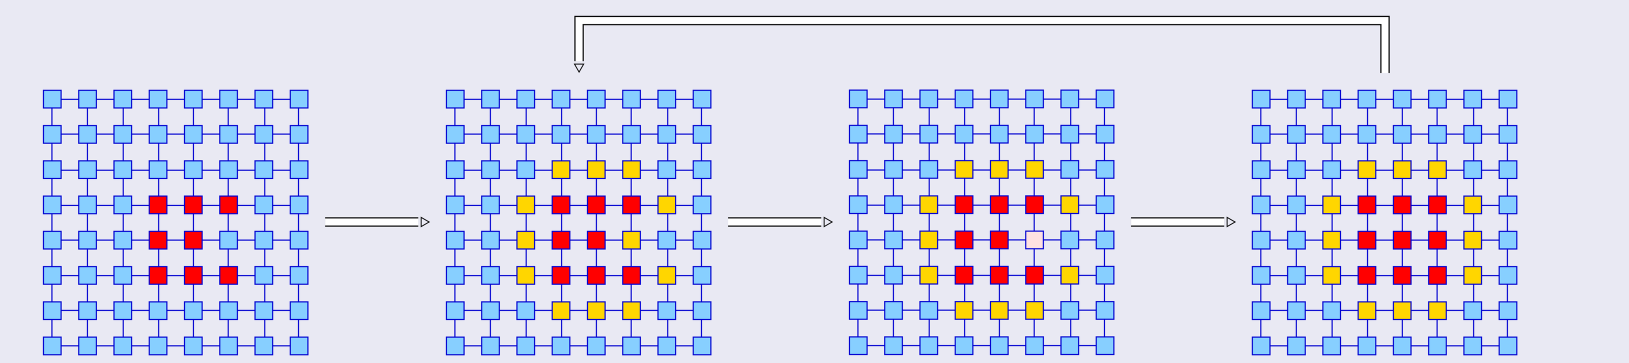

step 0
Solve LP relaxation
Reparametrize: $\theta \to \tilde{\theta}^{\phi}$

step 1
apply ILP solver to
+

step 2
check agreement
on the border

step 3
increase ILP subproblem
+ if disagree

Check the second column of this poster for definitions and analysis

Questions

- Is the consistency on the border () sufficient for optimality?
- How to select the initial LP/ILP (/) splitting?
- How to encourage consistency on the border ()?
- How to avoid re-solving the LP (+) part?
- Do we need to solve the LP (+) part to optimality?

See **Answers ...** below

Consistency on the Border is Insufficient for Optimality

Border () consistency alone is not enough. Example:

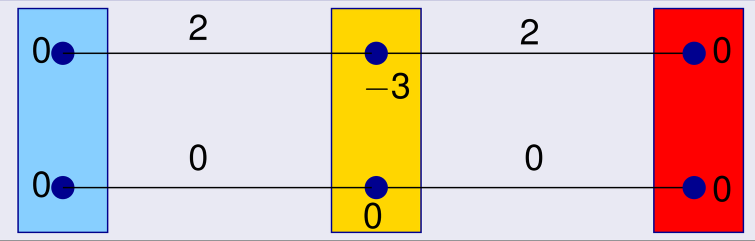

Background: Reparametrization

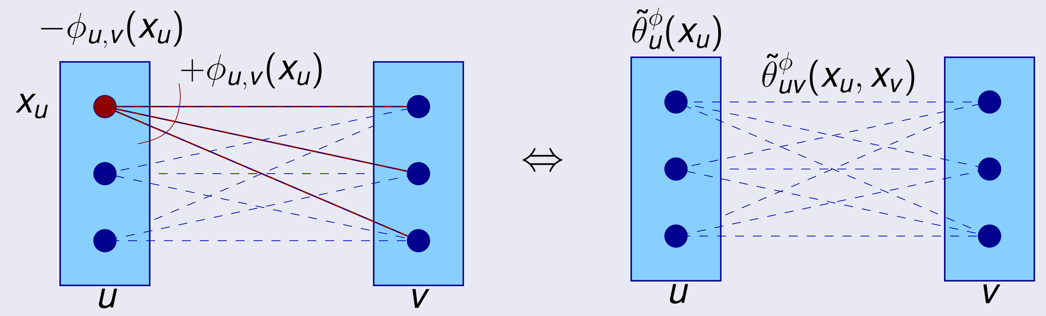

$$\text{Primal:} \quad E(\theta^{\phi}, x) = \min_x \sum_{v \in \mathcal{V}} \tilde{\theta}^{\phi}_v(x_v) + \sum_{uv \in \mathcal{E}} \tilde{\theta}^{\phi}_{uv}(x_u, x_v) = E(\theta, x)$$

$$\text{Dual:} \quad D(\phi) = \max_{\phi} \sum_{v \in \mathcal{V}} \underbrace{\min_{x_v} \tilde{\theta}^{\phi}_v(x_v)}_{\gamma_v} + \sum_{uv \in \mathcal{E}} \underbrace{\min_{x_{uv}} \tilde{\theta}^{\phi}_{uv}(x_{uv})}_{\gamma_{uv}} \leq E(\theta^{\phi}, x)$$

Background: Arc Consistency

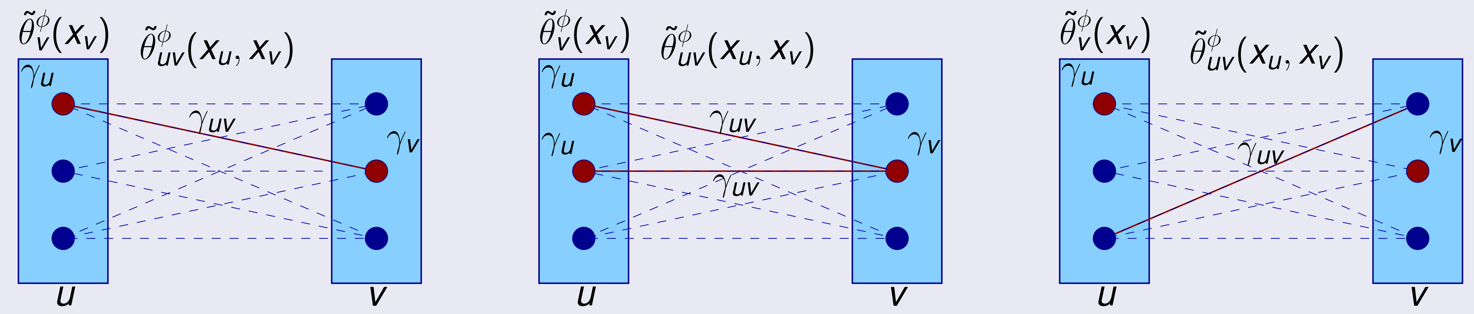

strict arc consistency — arc consistency — no arc consistency

Theorem. Strict arc consistency in all nodes ⇒ the *non-relaxed* problem is solved.

Main Theorem

Theorem. Let θ^{ϕ} be strictly arc consistent on +. Then if LP (+) and ILP (+) solutions agree on the border () their concatenation is globally optimal.

Answers: Why reparametrize?

Reparametrization provides:

- optimality condition (= *consistency on border* ())
- initial splitting criterion (*to and*)
- encouraging of border consistency
  Optimal labels "vote" for themselves in both LP (+) and ILP (+) subproblems
- potential speed-up of combinatorial solvers
  Acts as LP pre-solving

Moreover...

- An LP solver needs to be executed only *once*
  Because due to strict consistency local decisions are optimal: nodes removal does not change the remaining part of the solution
- *Suboptimal* reparametrization can be used as well
  Because we did not employ optimality of the reparametrization

Implementation Details

- Approximative LP solver → TRW-S [2]
- Arc consistency → tree agreement
- Combinatorial (ILP) solver → CPLEX

Results: Color Segmentation Potts models [1]

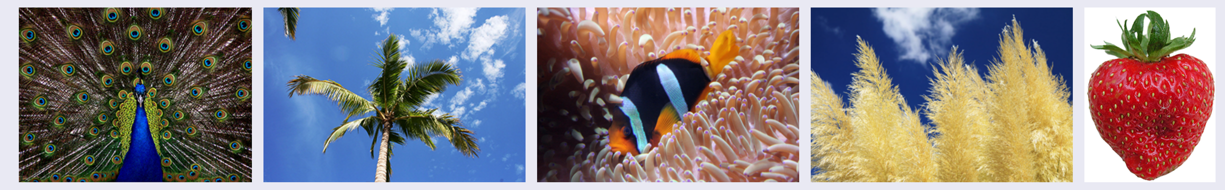

Overall 26 instances from `color segmentation` datasets 360 × 240, 4 − 12 labels and `brain` (up to $7 \cdot 10^6$ variables, 5 labels) – solved to optimality

Dataset	$E_{\mathcal{G},\theta}(x^*)$	LP step 0		ILP steps 1-3		MCA	MPLP		
		# it	time, s	# it	time, s	time, s	# LP it	LP time, s	ILP time, s
pfau	24010.44	1000	**276**	14	**14**	> 55496	10000	> 15000	
palm	12253.75	200	**65**	17	**93**	561	700	1579	3701
clownfish	14794.18	100	**32**	8	**10**	328	350	790	181
crops	11853.12	100	**32**	6	**6**	355	350	797	1601
strawberry	11766.34	100	**29**	8	**31**	483	350	697	1114

Table: Exemplary Potts model comparison on `Color segmentation (N8)` [1] dataset. Our method is **the fastest**.

Comparison to Partial Optimality Methods [3]

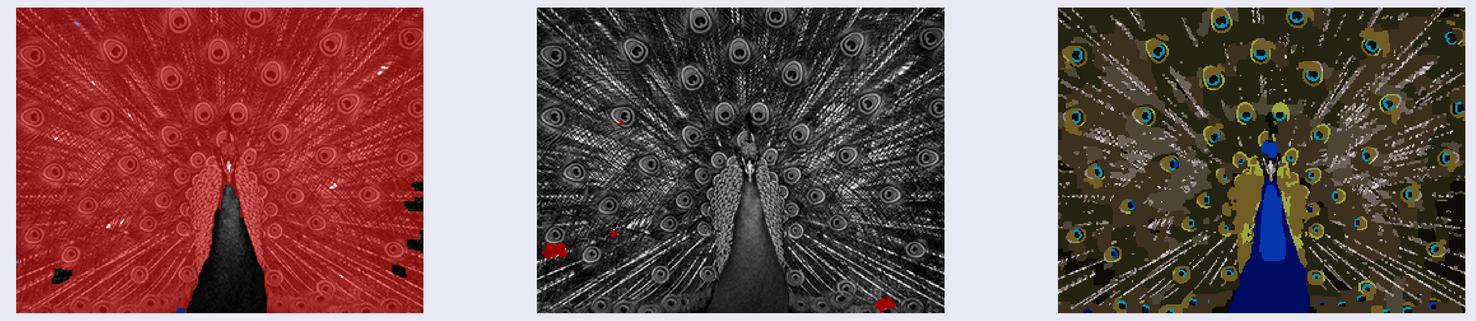

Method of Kovtun [3] — Our approach — Solution (remained unsolved so far!)

Figure: Red pixels mark nodes that need to be labeled by an ILP solver.

Results: Stereo [4, 1]

(a) tsukuba 384 × 288, 16 labels — solved
(b) venus 434 × 383, 20 labels — solved (remained unsolved so far!)
(c) teddy 450 × 375, 60 labels — not solved

Figure: Red pixels show the final subproblem passed to the ILP solver

Results: Photomontage [4, 1]

(a) family 752 × 566, 5 labels — solved (remained unsolved so far!)
(b) panorama 1071 × 480, 7 labels — not solved

Figure: Red points show the final subproblem passed to ILP solver (CPLEX)

References

[1] J. H. Kappes, B. Andres, F. A. Hamprecht, C. Schnörr, S. Nowozin, D. Batra, S. Kim, B. X. Kausler, J. Lellmann, N. Komodakis, and C. Rother. A comparative study of modern inference techniques for discrete energy minimization problems. In *CVPR*, 2013.

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

[3] I. Kovtun. Partial optimal labeling search for a NP-hard subclass of (max, +) problems. In *Proceedings of the DAGM Symposium*, 2003.

[4] 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 Trans. PAMI.*, 30:1068–1080, June 2008.

Acknowledgement. This work has been supported by the German Research Foundation (DFG) within the program "Spatio-/Temporal Graphical Models and Applications in Image Analysis", grant GRK 1653.