Partial Optimality by Pruning for MAP-inference with General Graphical Models

Paul Swoboda, Bogdan Savchynskyy, Jörg Kappes, Christoph Schnörr

Heidelberg University,
Germany

Segment the image...

Optimal labeling

Optimal labeling
NP hard

Optimal labeling
NP hard

- Solve convex relaxation (LP)
- Round relaxed solution
- NO optimality guarantees

Optimal labeling
NP hard

Partial labeling
Polynomially solvable
Optimality guaranteed

Energy Minimization - MAP-Inference with Graphical Models

$$\arg\min_{x\in X} J(x) := \sum_{f\in F} \theta_f(x_{ne(f)})$$

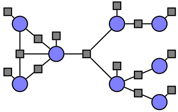

Factor graph
$G = (V, F, E)$

(circle)	- variable	$x_i \in \{1, \ldots, N\}$
(square)	- factor	$f \in F \subset V^n$
$\theta_f(x_{ne(f)})$	- potential of	$x_{ne(f)} \in \{1, \ldots, N\}^{\lvert ne(f) \rvert}$

Easy Examples

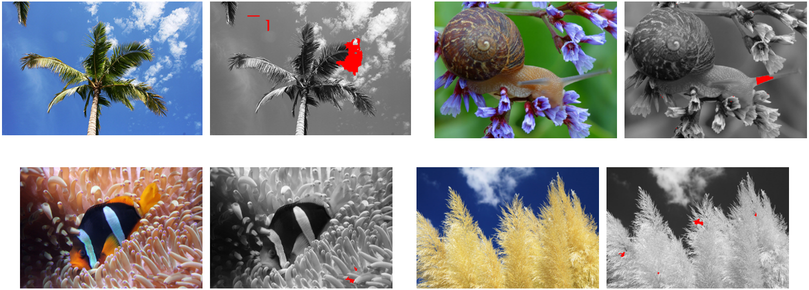

Figure : Color segmentation [Lellman 2010]

Difficult Examples

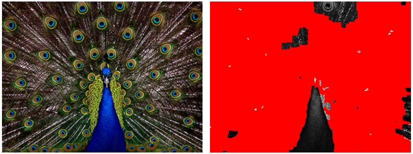

(a) color segmentation

(b) stereo

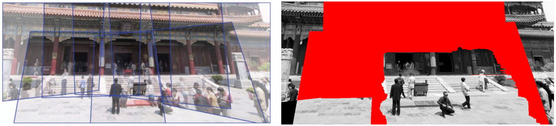

(c) panorama stitching

Figure : [Lellman 2010],[Szeliski et al. 2008],[Agarwala et al. 2004]

Related Work

Work	non-binary	higher order	non-Potts	Aux. problem
Boros & Hammer 2002	−	−	+	QPBO
Kovtun 2003	+	−	−	submodular
Rother et al. 2007	−	−	+	QPBO
Kohli et al. 2008	+	−	+	QPBO
Kovtun 2005	+	−	+	submodular
Fix et al. 2011	−	+	+	QPBO
Kahl & Strandmark 2012	−	+	+	bi-submodular
Windheuser et al. 2012	+	+	+	bi-submodular
Swoboda et al. 2013	+	−	−	LP
Shekhovtsov 2014	+	−	+	LP
Ours	+	+	+	**any relaxation**

Related Work

Work	non-binary	higher order	non-Potts	Aux. problem
Boros & Hammer 2002	−	−	+	QPBO
Kovtun 2003	+	−	−	submodular
Rother et al. 2007	−	−	+	QPBO
Kohli et al. 2008	+	−	+	QPBO
Kovtun 2005	+	−	+	submodular
Fix et al. 2011	−	+	+	QPBO
Kahl & Strandmark 2012	−	+	+	bi-submodular
Windheuser et al. 2012	+	+	+	bi-submodular
Swoboda et al. 2013	+	−	−	LP
Shekhovtsov 2014	+	−	+	LP
Ours	+	+	+	**any relaxation**

Related Work

Work	non-binary	higher order	non-Potts	Aux. problem
Boros & Hammer 2002	−	−	+	QPBO
Kovtun 2003	+	−	−	submodular
Rother et al. 2007	−	−	+	QPBO
Kohli et al. 2008	+	−	+	QPBO
Kovtun 2005	+	−	+	submodular
Fix et al. 2011	−	+	+	QPBO
Kahl & Strandmark 2012	−	+	+	bi-submodular
Windheuser et al. 2012	+	+	+	bi-submodular
Swoboda et al. 2013	+	−	−	LP
Shekhovtsov 2014	+	−	+	LP
Ours	+	+	+	**any relaxation**

Related Work

Work	non-binary	higher order	non-Potts	Aux. problem
Boros & Hammer 2002	−	−	+	QPBO
Kovtun 2003	+	−	−	submodular
Rother et al. 2007	−	−	+	QPBO
Kohli et al. 2008	+	−	+	QPBO
Kovtun 2005	+	−	+	submodular
Fix et al. 2011	−	+	+	QPBO
Kahl & Strandmark 2012	−	+	+	bi-submodular
Windheuser et al. 2012	+	+	+	bi-submodular
Swoboda et al. 2013	+	−	−	LP
Shekhovtsov 2014	+	−	+	LP
Ours	+	+	+	**any relaxation**

Algorithm Outline

Initialize: Generate labeling proposal
repeat
  Verify the proposal on a current graph
  Shrink the graph
until verification succeeds

Algorithm Outline

Initialize: Generate labeling proposal
repeat
  Verify the proposal on a current graph
  Shrink the graph
until verification succeeds

Algorithm Outline

Initialize: Generate labeling proposal
repeat
  Verify the proposal on a current graph
  Shrink the graph
until verification succeeds

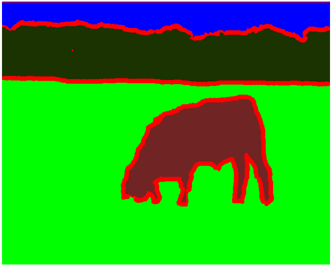

Algorithm Outline

Initialize: Generate labeling proposal
repeat
  Verify the proposal on a current graph
  Shrink the graph
until verification succeeds

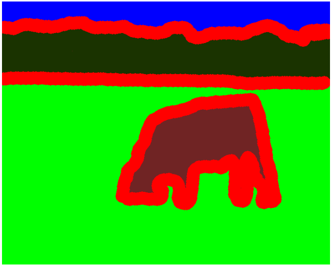

Algorithm

Proposed Partial Labeling

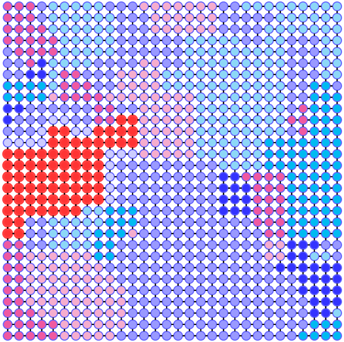

$$J(x) = \sum_{v \in V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u)$$

Labeling: $x \in X$

Partial optimality: $\chi \in \{0, 1\}^{|V|}$

Algorithm

Proposed Partial Labeling

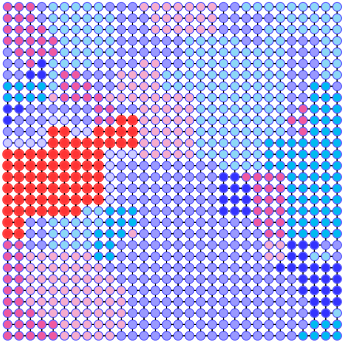

$$J(x) = \sum_{v \in V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u)$$

Labeling: $x \in X$

Partial optimality: $\chi \in \{0, 1\}^{|V|}$

Perturbed Problem

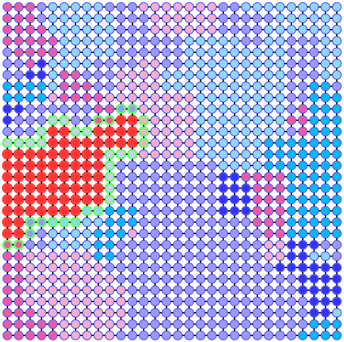

$$J_\chi(x) = \sum_{v \in V \cup V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u) + \sum_{v \in V} \bar{\theta}_v(x_v)$$

Algorithm

Proposed Partial Labeling

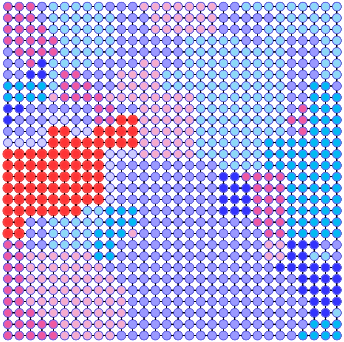

$$J(x) = \sum_{v \in V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u)$$

Labeling: $x \in X$

Partial optimality: $\chi \in \{0, 1\}^{|V|}$

Perturbed Problem

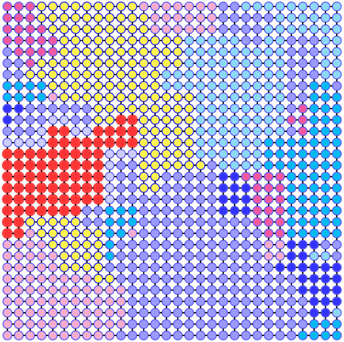

$$J_\chi(x) = \sum_{v \in V \cup V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u) + \sum_{v \in V} \bar{\theta}_v(x_v)$$

$\hat{x} = \arg\min_{x \in X} J_\chi(x)$

NP-hard → Relaxation

Algorithm

Proposed Partial Labeling

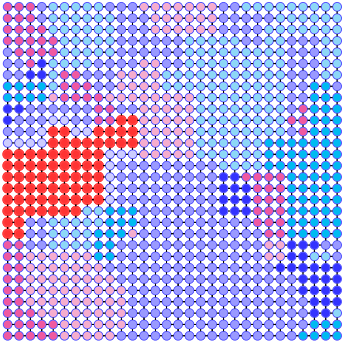

$$J(x) = \sum_{v \in V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u)$$

Labeling: $x \in X$

Partial optimality: $\chi \in \{0, 1\}^{|V|}$

Perturbed Problem

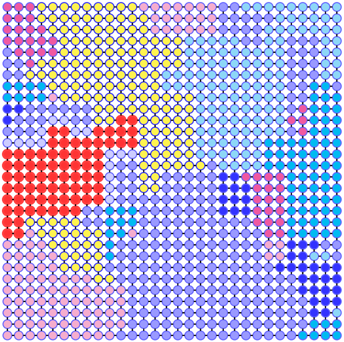

$$J_\chi(x) = \sum_{v \in V \cup V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u) + \sum_{v \in V} \bar{\theta}_v(x_v)$$

$\hat{x} = \arg\min_{x \in X} J_\chi(x)$ NP-hard → Relaxation

$\hat{x}_i \neq x_i \rightarrow \chi_i = 0$ Shrinking Rule

Algorithm

Proposed Partial Labeling

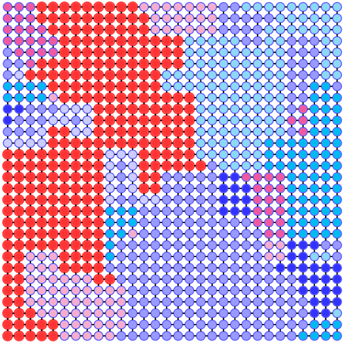

$$J(x) = \sum_{v \in V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u)$$

Labeling: $x \in X$

Partial optimality: $\chi \in \{0, 1\}^{|V|}$

Perturbed Problem

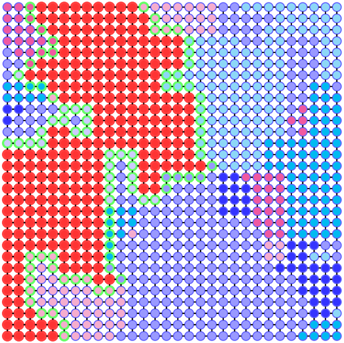

$$J_\chi(x) = \sum_{v \in V \cup V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u) + \sum_{v \in V} \bar{\theta}_v(x_v)$$

$\hat{x} = \arg\min_{x \in X} J_\chi(x)$ NP-hard → Relaxation

$\hat{x}_i \neq x_i \rightarrow \chi_i = 0$ Shrinking Rule

Algorithm

Proposed Partial Labeling

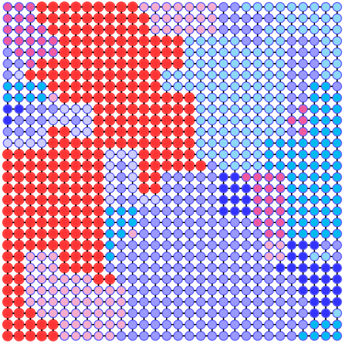

$$J(x) = \sum_{v \in V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u)$$

Labeling: $x \in X$

Partial optimality: $\chi \in \{0, 1\}^{|V|}$

Perturbed Problem

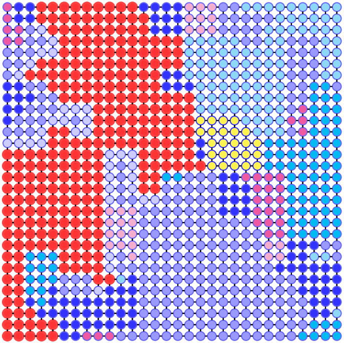

$$J_\chi(x) = \sum_{v \in V \cup V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u) + \sum_{v \in V} \bar{\theta}_v(x_v)$$

$\hat{x} = \arg\min_{x \in X} J_\chi(x)$ NP-hard → Relaxation

$\hat{x}_i \neq x_i \rightarrow \chi_i = 0$ Shrinking Rule

Algorithm

Proposed Partial Labeling

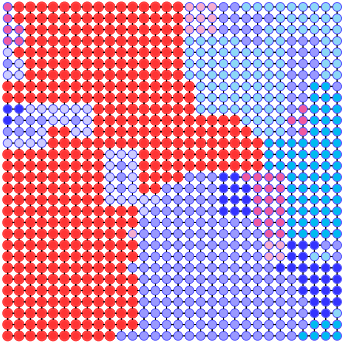

$$J(x) = \sum_{v \in V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u)$$

Labeling: $x \in X$

Partial optimality: $\chi \in \{0, 1\}^{|V|}$

Perturbed Problem

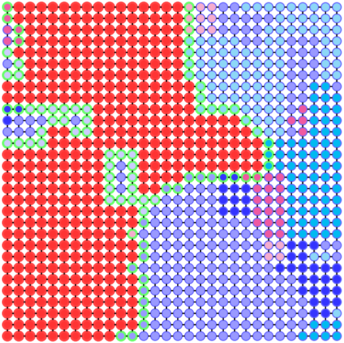

$$J_\chi(x) = \sum_{v \in V \cup V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u) + \sum_{v \in V} \bar{\theta}_v(x_v)$$

$\hat{x} = \arg\min_{x \in X} J_\chi(x)$ NP-hard → Relaxation

$\hat{x}_i \neq x_i \rightarrow \chi_i = 0$ Shrinking Rule

Algorithm

Proposed Partial Labeling

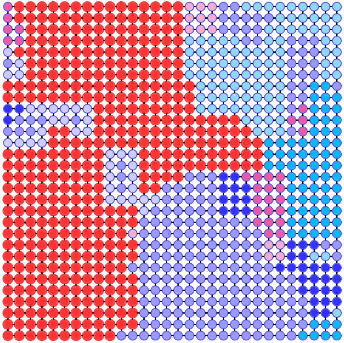

$$J(x) = \sum_{v \in V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u)$$

Labeling: $x \in X$

Partial optimality: $\chi \in \{0, 1\}^{|V|}$

Perturbed Problem

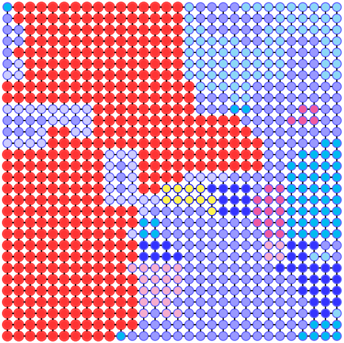

$$J_\chi(x) = \sum_{v \in V \cup V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u) + \sum_{v \in V} \bar{\theta}_v(x_v)$$

$\hat{x} = \arg\min_{x \in X} J_\chi(x)$ NP-hard → Relaxation

$\hat{x}_i \neq x_i \rightarrow \chi_i = 0$ Shrinking Rule

Algorithm

Proposed Partial Labeling

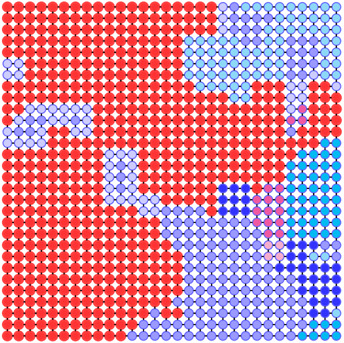

$$J(x) = \sum_{v \in V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u)$$

Labeling: $x \in X$

Partial optimality: $\chi \in \{0, 1\}^{|V|}$

Perturbed Problem

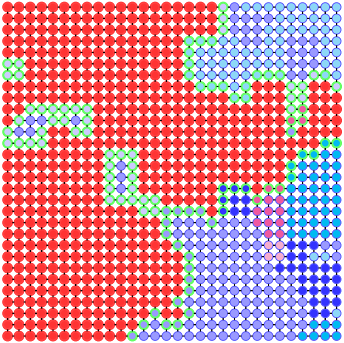

$$J_\chi(x) = \sum_{v \in V \cup V} \theta_v(x_v) + \sum_{uv \in E} \theta_{uv}(x_v, x_u) + \sum_{v \in V} \bar{\theta}_v(x_v)$$

$\hat{x} = \arg\min_{x \in X} J_\chi(x)$ NP-hard → Relaxation

$\hat{x}_i \neq x_i \rightarrow \chi_i = 0$ Shrinking Rule

Algorithm

Proposed Partial Labeling

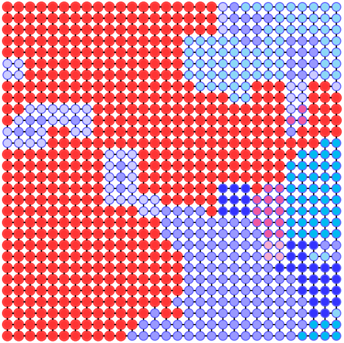

$$J(x) = \sum_{v\in V} \theta_v(x_v) + \sum_{uv\in E} \theta_{uv}(x_v, x_u)$$

Labeling: $x \in X$

Partial optimality: $\chi \in \{0, 1\}^{|V|}$

Perturbed Problem

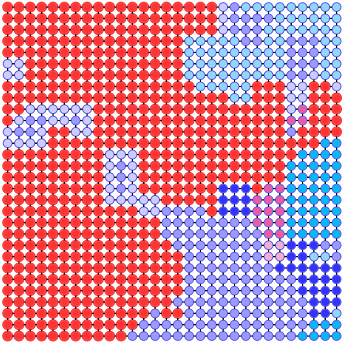

$$J_\chi(x) = \sum_{v\in V\cup V} \theta_v(x_v) + \sum_{uv\in E} \theta_{uv}(x_v, x_u) + \sum_{v\in V} \bar{\theta}_v(x_v)$$

$\hat{x} = \arg\min_{x\in X} J_\chi(x)$ NP-hard → Relaxation

$\hat{x}_i \neq x_i \rightarrow \chi_i = 0$ Shrinking Rule

Use of Approximate Solvers

Do we need to solve the relaxed problem

$$\hat{x} = \arg\min_{x \in X} J_{\chi}(x)$$

exactly?

NO!

Approximate solvers with optimality certificate
(like TRW-S [Kolmogorov 2005])
are allowed here

Results

Experiment (N)	MQPBO	Kovtun	GRD	Fix	Ours
`teddy`	0	†	†	†	**0.4423**
`venus`	0	†	†	†	**0.0009**
`family`	0.0432	†	†	†	**0.0611**
`pano`	0.1247	†	†	†	**0.5680**
`Potts (12)`	0.1839	0.7475	†	†	**0.9231**
`side-chain (21)`	0.0247	†	†	†	**0.6513**
`protein (8)`	†	†	0.2603	0.2545	**0.7799**
`cell-tracking`	†	†	†	0.1771	**0.9992**
`geo-surf (50)`	†	†	†	†	**0.8407**

Table : Percentage of persistent variables; † - method inapplicable.
We used local polytope relaxation and TRW-S and CPLEX as solvers.

Benchmarks: [Szeliski et al. 2008],[Kappes et al. 2013], [PIC 2011]

Results

Experiment (N)	MQPBO	Kovtun	GRD	Fix	Ours
teddy	0	†	†	†	**0.4423**
venus	0	†	†	†	**0.0009**
family	0.0432	†	†	†	**0.0611**
pano	0.1247	†	†	†	**0.5680**
Potts (12)	0.1839	0.7475	†	†	**0.9231**
side-chain (21)	0.0247	†	†	†	**0.6513**
protein (8)	†	†	0.2603	0.2545	**0.7799**
cell-tracking	†	†	†	0.1771	**0.9992**
geo-surf (50)	†	†	†	†	**0.8407**

Table : Percentage of persistent variables; † - method inapplicable.
We used local polytope relaxation and TRW-S and CPLEX as solvers.

Benchmarks: [Szeliski et al. 2008],[Kappes et al. 2013], [PIC 2011]

Potts models: Results

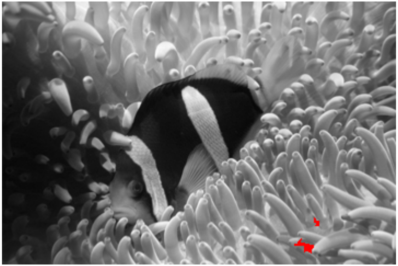

Ours

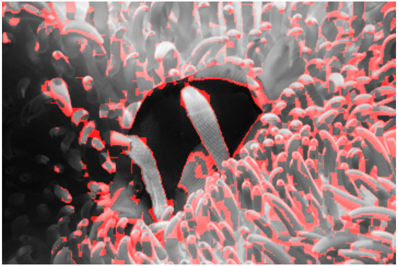

Kovtun's method

Potts models: Results

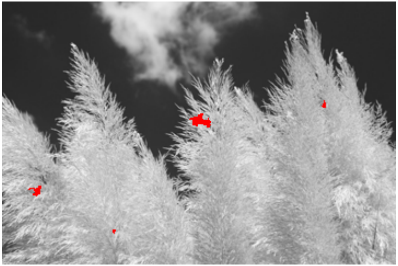

Ours

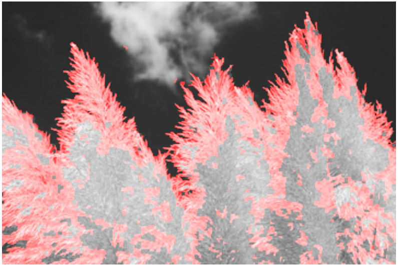

Kovtun's method

Potts models: Results

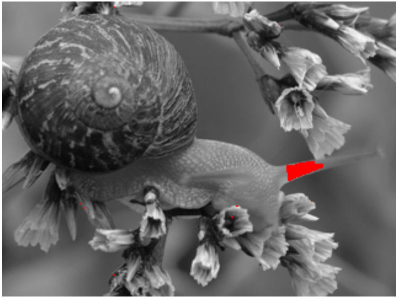

Ours

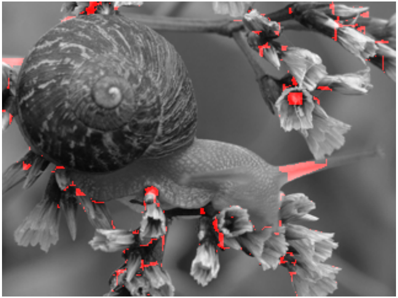

Kovtun's method

Potts models: Results

Ours

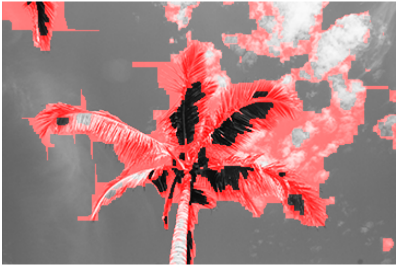

Kovtun's method

Potts models: Results

Ours

Kovtun’s method

Take Home Message and Outlook

We presented:

A generic method for partial optimality from MAP-Inference, which

- can employ any relaxation
- can use certain approximate solvers in the loop (e.g. TRW-S)
- scales as well as the used MAP-inference solver

Take Home Message and Outlook

We presented:

A generic method for partial optimality from MAP-Inference, which

- can employ any relaxation
- can use certain approximate solvers in the loop (e.g. TRW-S)
- scales as well as the used MAP-inference solver

Code

- preliminary research code at `http://paulswoboda.net`
- revised code will be included to OpenGM library soon.

Take Home Message and Outlook

We presented:

A generic method for partial optimality from MAP-Inference, which

- can employ any relaxation
- can use certain approximate solvers in the loop (e.g. TRW-S)
- scales as well as the used MAP-inference solver

Code

- preliminary research code at `http://paulswoboda.net`
- revised code will be included to OpenGM library soon.

Future work

- benefit on research on convex relaxations (tightness and speed)
- use finer partial optimality criteria (on label level)

Take Home Message and Outlook

We presented:

A generic method for partial optimality from MAP-Inference, which

- can employ any relaxation
- can use certain approximate solvers in the loop (e.g. TRW-S)
- scales as well as the used MAP-inference solver

Code

- preliminary research code at `http://paulswoboda.net`
- revised code will be included to OpenGM library soon.

Future work

- benefit on research on convex relaxations (tightness and speed)
- use finer partial optimality criteria (on label level)

Visit our poster number **O-2B-5**