

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

Paul Swoboda, Alexander Shekhovtsov, Jörg Hendrik Kappes,
Christoph Schnörr, and Bogdan Savchynskyy

Abstract—We consider the energy minimization problem for undirected graphical models, also known as MAP-inference problem for Markov random fields which is NP-hard in general. We propose a novel polynomial time algorithm to obtain a part of its optimal *non-relaxed integral* solution. Our algorithm is initialized with variables taking integral values in the solution of a convex relaxation of the MAP-inference problem and iteratively prunes those, which do not satisfy our criterion for partial optimality. We show that our pruning strategy is in a certain sense theoretically optimal. Also empirically our method outperforms previous approaches in terms of the number of persistently labelled variables. The method is very general, as it is applicable to models with arbitrary factors of an arbitrary order and can employ any solver for the considered relaxed problem. Our method's runtime is determined by the runtime of the convex relaxation solver for the MAP-inference problem.

Index Terms—MAP-inference, Markov random fields, energy minimization, persistency, partial optimality, local polytope

1 INTRODUCTION

FINDING the most likely configuration of a Markov random field (MRF), also called MAP-inference or energy minimization problem for graphical models, is of big importance in computer vision, bioinformatics, communication theory, statistical physics, combinatorial optimization, signal processing, information retrieval and statistical machine learning, see [1], [13], [39] for an overview of applications. This key problem however is NP-hard. Therefore approximate methods have been developed to tackle big instances commonly arising in image processing, see [13], [36] for an overview of such methods. These approximate methods often cannot find an optimal configuration, but deliver close solutions. If one could prove, that some variables of the solution given by such approximate algorithms belong to an optimal configuration, the value of such approximate methods would be greatly enhanced. In particular, the problem for the remaining variables could be solved by stronger, but computationally more expensive methods to obtain a global optimum as done, e.g., in [15].

In this paper we propose a way to gain such a partially optimal solution for the MAP-inference problem with *general* discrete MRFs from possibly also non-exact solutions of the commonly used local polytope relaxation (see [39]). Solving over the local polytope amounts to solving a linear problem for which *any* linear programming (LP) solver can be used and for which dedicated and efficient algorithms exist.

1.1 Related Work

We distinguish two classes of approaches to partial optimality.

- (i) *Roof duality based approaches.* The earliest paper dealing with persistency is [22], which states a persistency criterion for the stable set problem and verifies it for every solution of a certain relaxation. This relaxation is the same, as used by the roof duality method in [2] and which is also the basis for the well known QPBO-algorithm [2], [23]. The MQPBO method [17] extends roof duality to the multi-label case. The authors transform multi-label problems into quadratic binary ones and solve them via QPBO [2]. However, their transformation is dependent upon choosing a label order and their results are so as well, see the experiments in [34], where the label order is sampled randomly. It is not known how to choose an optimal label order to obtain the maximum number of persistent variables.

The roof duality method has been extended to higher order binary problems in [4], [10], [12]. The generalized roof duality (GRD) method for binary higher order problems [12] computes partially optimal variables directly for higher order potentials, while Ishikawa's and Fix et al's approaches [4], [10] transform the higher order problem to one with unary and pairwise terms only. Fix et al's method [4] is an improvement upon Ishikawa's [10].

- P. Swoboda and J.H. Kappes are with the Image and Pattern Analysis Group (IPA), Heidelberg University, Speyerer Str. 6, 69115 Heidelberg, Germany. E-mail: {swoboda, kappes}@math.uni-heidelberg.de.
- A. Shekhovtsov is with the Institute for Computer Graphics and Vision (ICG), Graz University of Technology, Inffeldgasse 16, Graz 8010, Austria. E-mail: shekhovtsov@icg.tugraz.at.
- C. Schnörr is with the Image and Pattern Analysis Group (IPA), Heidelberg University, Speyerer Str. 6, 69115 Heidelberg, Germany, and the Heidelberg Collaboratory for Image Processing (HCI), Heidelberg University, Speyerer Str. 6, 69115 Heidelberg, Germany. E-mail: schnoerr@math.uni-heidelberg.de.
- B. Savchynskyy is with the Heidelberg Collaboratory for Image Processing (HCI), Heidelberg University, Speyerer Str. 6, 69115 Heidelberg, Germany. E-mail: bogdan.savchynskyy@ivr.uni-heidelberg.de.

Manuscript received 14 Oct. 2014; revised 24 Mar. 2015; accepted 16 Sept. 2015. Date of publication 12 Oct. 2015; date of current version 10 June 2016.

Recommended for acceptance by A. Martinez, R. Basri, R. Vidal, and C. Fermüller.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPAMI.2015.2484327

TABLE 1
Comparison between Partial Optimality Methods

Work	Non-binary	Higher order	Non-Potts	Auxiliary problem
Boros & Hammer 2002 [2]	–	–	+	QPBO
Kovtun 2003[20]	+	–	–	submodular
Rother et al. 2007 [23]	–	–	+	QPBO
Kohli et al. 2008 [17]	+	–	+	QPBO
Kovtun 2011 [21]	+	–	+	submodular
Ishikawa 2011 [10]	–	+	+	QPBO
Fix et al. 2011 [4]	–	+	+	QPBO
Kahl & Strandmark 2012 [12]	–	+	+	bi-submodular
Windheuser et al. 2012 [40]	+	+	+	bi-submodular
Swoboda et al. 2013 [35]	+	–	–	local polytope
Shekhovtsov 2014 [29]	+	–	+	general linear program
Ours	+	+	+	any convex relaxation

Detailed description is presented in Section 1.1.

Windheuser et al. [40] proposed a *multi-label higher-order* roof duality method, which is a generalization of both MQPBO [17] to higher order and Kahl and Strandmark’s work [12] to the multi-label case. However Windheuser et al. neither describe an implementation nor provide experimental validation for the higher order multi-label case.

- (ii) *Labeling testing approaches.* A different approach, specialized for Potts models, is pursued by Kovtun [20], where possible labelings are tested for persistency by auxiliary submodular problems. The dead-end elimination procedure [3] tests, if certain labels of nodes cannot belong to an optimal solution. It is a local heuristic and does not perform any optimization.

Since for non-binary multi-labeling problems the submodular approximations constructed by approaches of class (i) are provably less tight than the standard local polytope relaxation [29, Prop. 1], we consider class (ii) in this paper. Specifically, based on ideas in [34] to handle the Potts model, we develop a theoretically substantiated approach to recognizing partial optimality for *general* graphical models, together with a competitive comparison to the five approaches [4], [10], [12], [17], [20] discussed above, that define the state-of-the-art.

Unified study. In addition we point to the recent paper [29], which provides a unified study of most mentioned methods and a systematic way of their analysis. While their persistency criterion is provably not weaker than ours, due to the general structure of the resulting LP it cannot be applied to large-scale problems in a straightforward manner. Moreover, our approach is directly applicable to higher order models and tighter than the local polytope relaxations, whereas [29] requires generalization to these cases, though such a generalization is presumably possible. We show that our algorithm solves a special case of the maximal persistency problem formulated in [29].

Shrinking technique. The recent work [25] proposes a method for efficient shrinking of the combinatorial search area with the local polytope relaxation. Though the algorithmic idea is similar to the presented one, the method [25] does not provide partially optimal solutions. We refer to Section 4 for further discussion.

Furthermore, preliminary shorter version of the our study was published at a conference as [35].

1.2 Contribution and Organization

Adopting ideas from [34], we propose a *novel* method for computing partial optimality, which is applicable to *general graphical models with arbitrary higher order potentials*. Similarly to [34] our algorithm is initialized with variables taking integral values in the solution of a convex relaxation of the MAP-inference problem and iteratively prunes those, which do not satisfy our persistency criterion. We show that our pruning strategy is in a certain sense theoretically optimal. Though the used relaxation can be chosen arbitrarily, for brevity we restrict our exposition and experiments to the local polytope relaxation. Tighter relaxations *provably* yield better results. However even by using the local polytope relaxation we can often achieve a *substantially higher* number of persistent variables, than competing approaches, which we confirm experimentally. We also show how our approach can be made invariant against reparametrizations. This improves our partial optimality criterion and we can show equivalence with the all-to-one improving mapping class of partial optimality methods proposed in [29]. Our approach is very general, as it can use *any*, also approximate, solver for the considered convex relaxation. Moreover, the computational complexity of our method is determined mainly by the runtime of the used solver.

The comparison to existing persistency methods is summarized in Table 1.

Our code together with the experimental setup is available at <http://paulswoboda.net>.

Organization. In Section 2 we review the energy minimization problem and the local polytope relaxation, in Section 3 our persistency criterion is presented. The corresponding algorithm and its theoretical analysis are presented in Sections 4, 5 and 6 respectively. Extensions to the higher order case and tighter relaxations are discussed in Section 7. Section 8 provides experimental validation of our approach and a comparison to the existing methods [4], [10], [12], [17], [20].

2 MAP-INFERENCE PROBLEM

The MAP-inference problem for a graphical model over an undirected graph $G = (\mathcal{V}, \mathcal{E})$, reads

$$\min_{x \in X_{\mathcal{V}}} E_{\mathcal{V}}(x) := \sum_{v \in \mathcal{V}} \theta_v(x_v) + \sum_{uv \in \mathcal{E}} \theta_{uv}(x_u, x_v), \quad (2.1)$$

where x_u belongs to a finite label set X_u for each node $u \in \mathcal{V}$, $\theta_u : X_u \rightarrow \mathbb{R}$ and $\theta_{uv} : X_u \times X_v \rightarrow \mathbb{R}$ are the *unary* and *pairwise potentials* associated with the nodes and edges of G . The label space for $A \subset \mathcal{V}$ is $X_A = \otimes_{u \in A} X_u$, where \otimes stands for the Cartesian product. For notational convenience we write $X_{uv} = X_u \times X_v$ and $x_{uv} = (x_u, x_v)$ for $uv \in \mathcal{E}$. Notations like $x \in X_A$ implicitly indicate that the vector x only has components x_u indexed by $u \in A$. With $x_{|A} \in X_A$ we denote restriction of the labeling $x \in X_{\mathcal{V}}$ to the set $A \subset \mathcal{V}$.

More general graphical models with terms depending on three or more variables can be considered as well. For brevity we restrict ourselves here to the pairwise case. An extension to the higher order case is discussed in Section 7.

Problem (2.1) is equivalent to the integer linear problem

$$\begin{aligned} \min_{\mu \in \Lambda_{\mathcal{V}}} \sum_{v \in \mathcal{V}} \sum_{x_v \in X_v} \theta_v(x_v) \mu_v(x_v) + \sum_{uv \in \mathcal{E}} \sum_{x_{uv} \in X_{uv}} \theta_{uv}(x_{uv}) \mu_{uv}(x_{uv}) \\ \text{s.t. } \mu_w(x_w) \in \{0, 1\} \text{ for } w \in \mathcal{V} \cup \mathcal{E}, x_w \in X_w, \end{aligned} \quad (2.2)$$

where the local polytope $\Lambda_{\mathcal{V}}$ [38] is the set of μ fulfilling

$$\begin{aligned} \sum_{x_v \in \mathcal{V}} \mu_v(x_v) &= 1, \quad v \in \mathcal{V}, \\ \sum_{x_v \in \mathcal{V}} \mu_{uv}(x_u, x_v) &= \mu_u(x_u), \quad x_u \in X_u, \quad uv \in \mathcal{E}, \\ \sum_{x_u \in \mathcal{V}} \mu_{uv}(x_u, x_v) &= \mu_v(x_v), \quad x_v \in X_v, \quad uv \in \mathcal{E}, \\ \mu_w(x_w) &\geq 0, \quad w \in \mathcal{V} \cup \mathcal{E}, \quad x_w \in X_w. \end{aligned} \quad (2.3)$$

We define Λ_A for $A \subset \mathcal{V}$ similarly. Slightly abusing notation we will denote the objective function in (2.2) as $E_{\mathcal{V}}(\mu)$. The formulation (2.2) utilizes the overcomplete representation [38] of labelings in terms of indicator vectors μ , which are often called *marginals*. The problem of finding $\mu^* \in \operatorname{argmin}_{\mu \in \Lambda_{\mathcal{V}}} E_{\mathcal{V}}(\mu)$ (i.e., solving (2.2) without integrality constraints) is called the *local polytope relaxation* of (2.1).

While solving the local polytope relaxation can be done in polynomial time, the corresponding optimal marginal μ^* may not be integral anymore, hence infeasible and not optimal for (2.2). For a wide spectrum of problems however most of the entries of optimal marginals μ^* for the local polytope relaxation will be integral. Unfortunately, there is no guarantee that any of these integral variables will be part of a globally optimal solution to (2.2), except in the case of binary variables, that is $X_u = \{0, 1\} \forall u \in \mathcal{V}$, and unary and pairwise potentials [6]. Natural questions are: (i) Is there a subset $A \subset \mathcal{V}$ and a minimizer μ^0 of the original NP-hard problem (2.2) such that $\mu_v^0 = \mu_v^* \forall v \in A$? In other words, is μ^* *partially optimal* or *persistent* on some set A ? (ii) Given a relaxed solution $\mu^* \in \Lambda_{\mathcal{V}}$, how can we determine such a set A ? We provide a novel approach to tackle these problems in what follows.

3 PERSISTENCY

Assume we have marginals $\mu \in \Lambda_{\mathcal{V}}$. We say that the marginal μ_u , $u \in \mathcal{V}$, is *integral* if $\mu_u(x_u) \in \{0, 1\} \forall x_u \in X_u$. In this case the marginal corresponds uniquely to a label x_u with $\mu_u(x_u) = 1$. If this *integrality condition* holds for all $u \in \mathcal{V}$ the corresponding vector μ will be denoted as $\delta(x)$. The convex

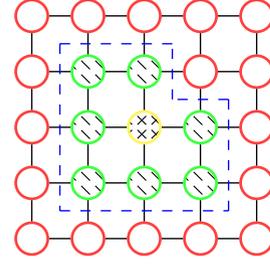


Fig. 1. An exemplary graph containing inside nodes (yellow with crosshatch pattern) and boundary nodes (green with diagonal pattern). The blue dashed line encloses the set A . Boundary edges are those crossed by the dashed line.

hull of marginals corresponding to all labelings known as *marginal polytope* will be denoted as $\mathcal{M}_{\mathcal{V}} := \operatorname{conv}(\delta(X_{\mathcal{V}}))$. The non-relaxed energy minimization (2.1) can be equivalently written as $\min_{\mu \in \mathcal{M}_{\mathcal{V}}} E_{\mathcal{V}}(\mu)$.

Let the boundary nodes and edges of a subset of nodes $A \subset \mathcal{V}$ be defined as follows:

Definition 1 (Boundary and interior). For the set $A \subset \mathcal{V}$ the set $\partial \mathcal{V}_A := \{u \in A : \exists v \in \mathcal{V} \setminus A \text{ s.t. } uv \in \mathcal{E}\}$ is called its *boundary*. The respective set of boundary edges is defined as $\partial \mathcal{E}_A = \{uv \in \mathcal{E} : u \in A \text{ and } v \in \mathcal{V} \setminus A\}$. The set $A \setminus \partial \mathcal{V}_A$ is called the *interior* of A .

An exemplary graph illustrating the concept of interior and boundary nodes can be seen in Fig. 1.

Definition 2 (Persistency). A labeling $x^0 \in X_A$ on a subset $A \subset \mathcal{V}$ is *partially optimal* or *persistent* if x^0 coincides with an optimal solution to (2.1) on A .

In the remainder of this section, we state our novel persistency criterion in Theorem 1. Taking additionally into account convex relaxation yields a computationally tractable approach in Corollary 1.

As a starting point, consider the following sufficient criterion for persistency of $x^0 \in X_A$. Introducing a *concatenation* of labelings $x^0 \in X_A$ and $\tilde{x} \in X_{\mathcal{V} \setminus A}$ as $(x^0, \tilde{x}) := \begin{cases} x_v^0, & v \in A, \\ \tilde{x}_v, & v \in \mathcal{V} \setminus A \end{cases}$, the criterion reads:

Proposition 1. The partial labeling $x^0 \in X_A$ is persistent if there holds

$$\forall \tilde{x} \in X_{\mathcal{V} \setminus A} : x^0 \in \operatorname{argmin}_{x \in X_A} E_{\mathcal{V}}((x, \tilde{x})). \quad (3.1)$$

Proof. Consider the equation

$$\min_{x \in X_{\mathcal{V}}} E(x) = \min_{\tilde{x} \in X_{\mathcal{V} \setminus A}} \min_{x \in X_A} E_{\mathcal{V}}((x, \tilde{x})). \quad (3.2)$$

Let $\tilde{x} \in X_{\mathcal{V} \setminus A}$ be such that it leads to a minimal value on the right hand side of (3.2). Then \tilde{x} is part of an optimal solution. By the assumption (3.1), x^0 is an optimal solution to the inner minimization problem of (3.2), hence (x^0, \tilde{x}) is optimal for (2.1). \square

This means that if we fix *any* labeling \tilde{x} on the complement of A and optimize with respect to x^0 on A , the concatenated labeling (x^0, \tilde{x}) has to be optimal. Informally this means that the solution x^0 is independent of what

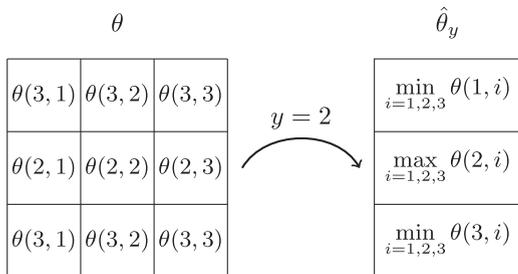


Fig. 2. Illustration of a boundary potential $\hat{\theta}_y$ constructed in (3.3). The second label comes from the test labeling y , therefore entries are maximized for the second row and minimized otherwise.

happens on $\mathcal{V} \setminus A$. This criterion however is hard to check directly, as it entails solving NP-hard minimization problems over an exponential number of labelings $\tilde{x} \in X_{\mathcal{V} \setminus A}$.

We relax the above criterion (3.1) so that we have to check the solution of only *one* energy minimization problem by modifying the unaries θ_v on boundary nodes so that they bound the influence of *all* labelings on $\mathcal{V} \setminus A$ uniformly.

Definition 3 (Boundary potentials and energies). For a set $A \subset \mathcal{V}$ and a test labeling $y \in X_A$, we define for each boundary edge $uv \in \partial\mathcal{E}_A$, $u \in \partial\mathcal{V}_A$ the “boundary” potential $\hat{\theta}_{uv,y_u} : X_u \rightarrow \mathbb{R}$ as follows:

$$\hat{\theta}_{uv,y_u}(x_u) := \begin{cases} \max_{x_v \in X_v} \theta_{uv}(x_u, x_v), & y_u = x_u \\ \min_{x_v \in X_v} \theta_{uv}(x_u, x_v), & y_u \neq x_u. \end{cases} \quad (3.3)$$

Define the energy $\hat{E}_{A,y} : X_A \rightarrow \mathbb{R}$ with test labeling y as

$$\hat{E}_{A,y}(x) := E_A(x) + \sum_{uv \in \partial\mathcal{E}_A : u \in \partial\mathcal{V}_A} \hat{\theta}_{uv,y_u}(x_u), \quad (3.4)$$

where $E_A(x) = \sum_{u \in A} \theta_u(x_u) + \sum_{uv \in \mathcal{E} : u,v \in A} \theta_{uv}(x_{uv})$ is the energy with potentials with support in A .

Given a test labeling $y \in X_A$, energy (3.4) assigns a higher value than the original energy (2.1) for all labelings conforming to y and makes it more favourable for all labelings to not conform to y . An illustration of a boundary potential is depicted by Fig. 2.

As a consequence, if the test labeling y from Definition 1 minimizes the energy (3.4), the proof of the following theorem asserts that changing an arbitrary labeling $x \in X_{\mathcal{V}}$ as follows: $x'(v) = \begin{cases} y(v), & v \in A \\ x(v), & v \notin A \end{cases}$ will always result in a labeling with not bigger energy (2.1), hence y in particular fulfills the conditions (3.1) of Proposition 1 and thus is persistent.

Theorem 1 (Partial optimality criterion). A labeling $x^0 \in X_A$ on a subset $A \subseteq \mathcal{V}$ is persistent if

$$x^0 \in \operatorname{argmin}_{x \in X_A} \hat{E}_{A,x^0}(x), \quad (3.5)$$

where \hat{E}_{A,x^0} is the augmented energy functional (3.4).

To prove the theorem we need the following technical lemma.

Lemma 1. Let $A \subset \mathcal{V}$ be given together with $y \in X_A$. Let x^0 and x' be two labelings on \mathcal{V} such that $x^0|_A = y$. Then it holds for

$uv \in \partial\mathcal{E}_A$, $u \in \partial\mathcal{V}_A$ that

$$\theta_{uv}(x_u^0, x_v') + \hat{\theta}_{uv,y}(x_u') - \hat{\theta}_{uv,y}(x_u^0) \leq \theta_{uv}(x_u', x_v'). \quad (3.6)$$

Proof. The case $x_u' = x_u^0$ is trivial. Otherwise, by Definition 3, inequality (3.6) is equivalent to

$$\theta_{uv}(x_u^0, x_v') + \min_{x_v \in X_v} \theta_{uv}(x_u', x_v) - \max_{x_v \in X_v} \theta_{uv}(x_u^0, x_v) - \theta_{uv}(x_u', x_v') \leq 0. \quad (3.7)$$

Choose x_v' for x_v in the minimization and maximization in (3.7) to obtain the result. \square

Proof of Theorem 1. Let

$$\tilde{x} \in \operatorname{arg} \min_{\substack{x \in X_{\mathcal{V}} \\ x|_A = x^0|_A}} E_{\mathcal{V}}(x), \quad (3.8)$$

and let $x' \in X_{\mathcal{V}}$ be an arbitrary labeling. Then

$$\begin{aligned} E_{\mathcal{V}}(\tilde{x}) &= E_A(x^0) + E_{\mathcal{V} \setminus A}(\tilde{x}) + \sum_{uv \in \partial\mathcal{E}_A} \theta_{uv}(x_u^0, \tilde{x}_v) \\ &= E_A(x^0) + \sum_{uv \in \partial\mathcal{E}_A} \hat{\theta}_{uv,y}(x_u^0) \\ &\quad + E_{\mathcal{V} \setminus A}(\tilde{x}) + \sum_{uv \in \partial\mathcal{E}_A} [\theta_{uv}(x_u^0, \tilde{x}_v) - \hat{\theta}_{uv,y}(x_u^0)] \\ &= \hat{E}_{A,x^0}(x^0) + E_{\mathcal{V} \setminus A}(\tilde{x}) + \sum_{uv \in \partial\mathcal{E}_A} [\theta_{uv}(x_u^0, \tilde{x}_v) - \hat{\theta}_{uv,x^0}(x_u^0)] \\ &\leq \hat{E}_{A,x^0}(x') + E_{\mathcal{V} \setminus A}(x') + \sum_{uv \in \partial\mathcal{E}_A} [\theta_{uv}(x_u^0, x'_v) - \hat{\theta}_{uv,x^0}(x_u^0)] \end{aligned} \quad (3.10)$$

$$\begin{aligned} &= E_A(x') + \sum_{uv \in \partial\mathcal{E}_A} \hat{\theta}_{uv,x^0}(x_u^0) \\ &\quad + E_{\mathcal{V} \setminus A}(x') + \sum_{uv \in \partial\mathcal{E}_A} [\theta_{uv}(x_u^0, x'_v) - \hat{\theta}_{uv,x^0}(x_u^0)] \\ &\leq E_A(x') + E_{\mathcal{V} \setminus A}(x') + \sum_{uv \in \partial\mathcal{E}_A} \theta_{uv}(x_u', x'_v) = E_{\mathcal{V}}(x'). \end{aligned} \quad (3.11)$$

The equality (3.9) is due to definition of \tilde{x} in (3.7). The first inequality 3.10 is due to $x^0 \in \operatorname{argmin}_x \hat{E}_{A,x^0}(x)$, as assumed, and of \tilde{x} for (3.8). The second inequality 3.11 is due to Lemma 1. Hence x^0 is part of a globally optimal solution, as x' was arbitrary. \square

Checking the criterion in Theorem 1 is NP-hard, because (3.5) is a MAP-inference problem of the same class as (2.1). By relaxing the minimization problem (3.5) one obtains the polynomially verifiable persistency criterion in Corollary 1.

Corollary 1 (Tractable partial optimality criterion). Labeling $x^0 \in X_A$ on $A \subset \mathcal{V}$ fulfilling the condition

$$\delta(x^0) \in \operatorname{argmin}_{\mu \in \Lambda_A} \hat{E}_{A,x^0}(\mu) \quad (3.12)$$

is also a solution to (3.5), hence persistent on A .

Proof. Expression (3.12) implies

$$\delta(x^0) \in \operatorname{argmin}_{\mu \in \Lambda_A, \mu \in \{0,1\}} \hat{E}_{A,x^0}(\mu) \quad (3.13)$$

because $\delta(x^0)$ is integral by definition. As (2.1) and (2.2) are equivalent and the corresponding labeling x^0 satisfies the conditions of Theorem 1, x^0 is partially optimal on A . \square

4 PERSISTENCY ALGORITHM

Now we concentrate on finding a set A and labeling $x \in X_A$ such that the solution of $\min_{\mu \in \Lambda_A} \hat{E}_{A,x}(\mu)$ fulfills the conditions of Corollary 1. Our approach is summarized in Algorithm 1.

In the initialization step of Algorithm 1 we solve the relaxed problem over \mathcal{V} without boundary labeling and initialize the set A^0 with nodes having an integer label. Then in each iteration t we minimize over the local polytope the energy \hat{E}_{A^t, x^t} defined in (3.4), corresponding to the set A^t and boundary labeling coming from the solution of the last iteration. We remove from A^t all variables which are not integral or do not conform to the boundary labeling. In each iteration t of Algorithm 1 we shrink the set A^t by removing variables taking non-integral values or not conforming to the current boundary condition. See Fig. 3 for an illustration of one iteration of Algorithm 1.

Algorithm 1. Finding Persistent Variables.

Data: $G = (\mathcal{V}, \mathcal{E}), \theta_u : X_u \rightarrow \mathbb{R}, \theta_{uv} : X_{uv} \rightarrow \mathbb{R}$

Result: $A^* \subset \mathcal{V}, x^* \in X_{A^*}$

Initialize:

Choose $\mu^0 \in \operatorname{argmin}_{\mu \in \Lambda_{\mathcal{V}}} E_{\mathcal{V}}(\mu)$

$A^0 = \{u \in \mathcal{V} : \mu_u^0 \in \{0, 1\}^{|X_u|}\}$

$t = 0$

repeat

Set x_u^t such that $\mu_u^t(x_u^t) = 1, u \in A^t$

Choose $\mu^{t+1} \in \operatorname{argmin}_{\mu \in \Lambda_{A^t}} \hat{E}_{A^t, x^t}(\mu)$

$t = t + 1$

$W^t = \{u \in \partial \mathcal{V}_{A^{t-1}} : \mu_u^t(x_u^{t-1}) \neq 1\}$

$A^t = \{u \in A^{t-1} : \mu_u^t \in \{0, 1\}^{|X_u|}\} \setminus W^t$

until $A^t = A^{t-1}$

$A^* = A^t$

Set $x^* \in X_{A^*}$ such that $\mu_u^t(x_u^*) = 1$

Convergence. Since \mathcal{V} is finite and $|A^t|$ is monotonically decreasing, the algorithm converges in at most $|\mathcal{V}|$ steps. Solving each subproblem in Algorithm 1 can be done in polynomial time. As the number of iterations of Algorithm 1 is at most $|\mathcal{V}|$, Algorithm 1 itself is polynomial as well. In practice only few iterations are needed.

After termination of Algorithm 1, we have

$$\delta(x^*) \in \operatorname{argmin}_{\mu \in \Lambda_{A^*}} \hat{E}_{A^*, x^*}(\mu). \quad (4.1)$$

Hence x^* and A^* fulfill the conditions of Corollary 1, which proves persistency.

Choice of solver. All our results are independent of the specific algorithm one uses to solve the relaxed problems $\min_{\mu \in \Lambda_A} \hat{E}_{A,y}$, provided it returns an exact solution. However this can be an issue for large-scale datasets, where classical exact LP solvers like, e.g., the simplex method become inapplicable. It is important that one can also employ approximate solvers, as soon as they provide (i) a proposal for potentially persistent nodes and (ii) sufficient conditions

for optimality of the found *integral* solutions such as, e.g., zero duality gap. These properties have the following precise formulation.

Definition 4 (Consistent labeling). A labeling $c \in \otimes_{v \in \mathcal{V}} (X_v \cup \{\#\})$ is called a consistent labeling for the energy minimization problem (2.1), if from $c_v \in X_v \forall v \in \mathcal{V}$ follows that $c \in \operatorname{argmin}_{x \in X_{\mathcal{V}}} E_{\mathcal{V}}(x)$.

We will call an algorithm for solving the energy minimization problem (2.1) consistency ascertaining, if it provides a consistent labeling as its output.

Consistent labelings can be constructed for a wide range of algorithms, e.g.:

- Dual decomposition based algorithms [14], [18], [19], [24], [26] deliver *strong tree agreement* [37] and algorithms considering the Lagrangian dual [5], [7], [28] return *strong arc consistency* [39] for some nodes. If one of these properties holds for a node v , we set c_v as the corresponding label. Otherwise we set $c_v = \#$.
- Naturally, any algorithm solving $\min_{\mu \in \Lambda_{\mathcal{V}}} E(\mu)$ exactly is consistency ascertaining with

$$c_v = \begin{cases} x_v, & \mu_v(x_v) = 1 \\ \#, & \mu_v \notin \{0, 1\}^{|X_v|}. \end{cases}$$

Proposition 2. Let operations $\mu \in \operatorname{argmin}(\dots)$ in Algorithm 1 be exchanged with

$$\forall v \in \mathcal{V}, x_v \in X_v, \mu_v(x_v) := \begin{cases} 1, & c_v = x_v \\ 0, & c_v \notin \{x_v, \#\}, \\ 1/|X_v|, & c_v = \# \end{cases}$$

where c are consistent labelings returned by a consistency ascertaining algorithm applied to the corresponding minimization problems. Then the output labeling x^* is persistent.

Proof. At termination of Algorithm 1 we have obtained a subset of nodes A^* , a test labeling $y^* \in X_{A^*}$, a labeling x^* equal to y^* on A and a consistency mapping $c_u = x_u^*$ for $u \in A^*$. Hence, by Definition 4, $x^* \in \operatorname{argmin}_{x \in X_{A^*}} \hat{E}_{A^*, y^*}$ and x^* fulfills the conditions of Theorem 1. \square

Remark 1. Note that a bad or early stopped solver, i.e., one which rarely (or even never) returns an optimality certificate or solves a weak relaxation, will also work with Algorithm 1. However it will find smaller (or even empty) partial optimal solutions.

Comparison to the Shrinking Technique (CombiLP) [25]. The recently published approach [25], similar to Algorithm 1, describes how to shrink the combinatorial search area with the local polytope relaxation. However (i) Algorithm 1 solves a series of auxiliary problems on the subsets A^t of integer labels, whereas the method [25] considers nodes, which got fractional labels in the relaxed solution; (ii) Algorithm 1 is polynomial and provides only persistent labels, whereas the method [25] has exponential complexity and either finds an optimal solution or gives no information about persistence.

From the practical point of view, both algorithms have different application scenarios: CombiLP [25] will only work on sparse graphs, as otherwise the combinatorial part, which one has to solve with exact methods, becomes too

TABLE 2
Comparison between Our Method and CombiLP [25]

	CombiLP [25]	Our method
Dense graphs	–	+
Very large-scale	–	+
Big fractional part of LP solution	–	+
Relaxed MAP-inference is solved only once	+	–
Provides a complete solution to Labeling Problem (2.1)	+	–

big, as the boundary $\partial\mathcal{V}_A$ for $A \subsetneq \mathcal{V}$ grows very quickly then. Also, even for sparse graphs, the combinatorial part may not grow too big during the application of the algorithm, as otherwise the combinatorial solver will again not be able to cope with it. Our algorithm does not possess these two disadvantages. From the perspective of running time it does not matter how big the set $\mathcal{V} \setminus A^t$ becomes during the iterations of Algorithm 1. On the other hand, the subsets of variables to which the method [25] applies a combinatorial solver to achieve global optimality are often smaller than $\mathcal{V} \setminus A^t$ in Algorithm 1, because potentials in CombiLP [25] remain unchanged in contrast to the perturbation (3.4). Another advantage of the method [25] is that it needs to solve the (typically) big LP relaxation of the original problem only once, whereas our method does this iteratively, which makes it often slower than CombiLP.

One other possible application scenario which is possible with our method but not with CombiLP [25] is the following: Assume we want to solve an extremely big inference problem, one that does not fit even into memory. To do this, choose a subset $A \subsetneq \mathcal{V}$ of nodes of the graphical model, solve the inference problem on the induced subgraph $G(A)$ with some boundary conditions, and find a partially optimal labeling on it. This is akin to the windowing technique of [29]. By doing so for an overlapping set of subgraphs, one may try to find a labeling for the overall problem on G .

The major differences between CombiLP [25] and our method are summarised in Table 2.

5 LARGEST PERSISTENT LABELING

Let $A^0 \subseteq \mathcal{V}$ and $\mu^0 \in \Lambda_{A^0}$ be defined as in Algorithm 1. Subsets $A \subset A^0$ which fulfill the conditions of Corollary 1 taken with labelings $\mu^0|_A$ can be partially ordered with respect to inclusion \subset of their domains. In this section we will show that the following holds:

- There is a largest set among those, for which there exists a *unique* persistent labeling fulfilling the conditions of Corollary 1.
- Algorithm 1 finds this largest set.

This will imply that Algorithm 1 cannot be improved upon with regard to the criterion in Corollary 1.

Definition 5 (Strong persistency). A labeling $x^* \in X_A$ is called strongly persistent on A , if x^* is the unique labeling on A fulfilling the conditions of Theorem 1.

Lemma 2. Let $x^* \in X_A$ be strongly persistent. Then for any optimal solution x of (2.1) we have $x^* = x|_A$.

Proof. This follows from Inequality 3.10 being strict in this case. \square

Theorem 2 (Largest persistent labeling). Algorithm 1 finds a superset A^* of the largest set $A^{*strong} \subseteq A^* \subset \mathcal{V}$ of strongly persistent variables identifiable by the criterion in Corollary 1.

To prove the theorem we need the following technical lemma.

Lemma 3. Let $A \subset B \subset \mathcal{V}$ be two subsets of \mathcal{V} and $\mu^A \in \Lambda_A$ marginals on A and $x^A \in X_A$ a labeling fulfilling the conditions of Corollary 1 uniquely (i.e., x^A is strongly persistent). Let $y^B \in X_B$ be a test labeling such that $y^B|_A = x^A$.

Then for all marginals $\mu^* \in \operatorname{argmin}_{\mu \in \Lambda_B} \hat{E}_{B,y^B}(\mu)$ on B it holds that $\mu_v^*(x_v^A) = 1 \forall v \in A$.

Proof. Similar to the proof of Theorem 1. Replace \mathcal{V} by B . \square

Proof of Theorem 2 We will use the notation from Algorithm 1. It will be enough to show that for every $\bar{A} \subseteq \mathcal{V}$ such that there exists a strongly persistent labeling $\bar{x} \in X_{\bar{A}}$ we have $\bar{A} \subseteq A^t$ in each iteration of Algorithm 1 and furthermore $\bar{x}_v = x_v^t$ for all $v \in \mathcal{V}_{\bar{A}}$. Hence the union of sets $A^{tstrong}$ for which a strongly persistent labeling exists which fulfills the conditions of Corollary 1, is a subset of $A^t \forall t$. Also by Lemma 2 the associated strongly persistent labelings agree where they overlap, hence we are done.

For $t=0$ apply Lemma 3 with $A := \bar{A}$ and $B := A^0 (= \mathcal{V})$. Condition $\bar{x} = y^B|_{\bar{A}}$ in Lemma 3 is assured by Corollary 1. Hence, Lemma 3 ensures that for all $\mu^0 \in \operatorname{argmin}_{\mu \in \Lambda_{\mathcal{V}}} E(\mu)$ it holds that $\mu_v^0(\bar{x}_v) = 1$ for all $v \in \bar{A}$.

Now assume the claim to hold for iteration $t-1$. We need to show that it also holds for t . For this invoke Lemma 3 with $A := \bar{A}$, $B := A^{t-1}$ and $y^B := x^{t-1}$. The conditions of Lemma 3 hold by assumption on $t-1$. Lemma 3 now ensures that for all $\mu^t \in \operatorname{argmin}_{\mu \in \Lambda_{A^{t-1}}} \hat{E}_{A^{t-1},x^{t-1}}(\mu)$ there holds $\mu^t(x_v^A) = 1 \forall v \in A$. \square

From the proof of Theorem 2 we can directly conclude the existence and uniqueness of a largest strongly persistent labeling identifiable by Corollary 1 and a set supporting it.

Corollary 2. There exists a unique largest set $A^{*strong}$, for which there exists a strongly persistent labeling identifiable by Corollary 1.

Also exactly the largest strongly persistent labeling identifiable by Corollary 1 can be found under a mild uniqueness assumption.

Corollary 3. If there is a unique solution of $\min_{\mu \in \Lambda_{A^t}} \hat{E}_{A^t,x^t}(\mu)$ for all $t = 0, \dots$ obtained during the iterations of Algorithm 1, then Algorithm 1 finds the largest subset of persistent variables identifiable by the sufficient partial optimality criterion in Corollary 1.

Remark 2. Above we showed that Algorithm 1 will find a persistent labeling which contains the largest strongly persistent one identifiable by Corollary 1. The two may differ when the optimization problems solved in the course of Algorithm 1 have multiple optima. The simplest example of such a situation occurs if the relaxation

$\min_{\mu \in \Lambda_{\mathcal{V}}} E_{\mathcal{V}}(\mu)$ is tight, but has several integer solutions. Any convex combination of these solutions will form a non-integral solution, hence the strongly persistent labeling is defined on a smaller set than any integral solution of $\min_{\mu \in \Lambda_{\mathcal{V}}} E_{\mathcal{V}}(\mu)$, which is non strongly persistent. Note however that a labeling obtained by Algorithm 1, also when it is not strongly persistent, comes from *one* globally optimal labeling, i.e., it can be completed to a globally optimal labeling by solving for the remaining variables.

6 REPARAMETRIZATION AND OPTIMALITY OF THE METHOD

It is well-known [27] (see also [39]) that representation (2.1) of the energy function is not unique. There are other potentials, which keep the energy of all labelings unchanged. Any such potentials θ^{φ} can be represented as

$$\theta_v^{\varphi}(x_v) := \theta_v(x_v) - \sum_{u \in \text{nb}(v)} \varphi_{v,u}(x_v), \quad (6.1)$$

$$\theta_{uv}^{\varphi}(x_u, x_v) := \theta_{uv}(x_u, x_v) + \varphi_{v,u}(x_v) + \varphi_{u,v}(x_u) \quad (6.2)$$

with some numbers $\varphi_{u,v}(x_u)$, $uv \in \mathcal{E}$, $x_u \in X_u$, where $\text{nb}(v) := \{u \in \mathcal{V} : uv \in \mathcal{E}\}$ denotes the set of nodes adjacent to $v \in \mathcal{V}$. The vector φ with coordinates $\varphi_{u,v}(x_u)$ is called *reparametrization*.

6.1 Optimal Reparametrization

The boundary potentials (3.3) and hence the persistency approach described above are dependent on reparametrization. The natural question is existence of an *optimal reparametrization*, that is, the one providing the largest persistent set.

The only coordinates of the reparametrization vector φ , which can potentially influence the solution of the test problem (3.5) are $\varphi_{v,u}(x_v)$, $u \in \partial \mathcal{V}_A$, $uv \in \partial \mathcal{E}_A$. Reparametrization $\varphi_{v,u}(x_v)$, $v \in A$ “inside” A does not influence the solution, because it does not change the augmented energy $\hat{E}_{A,\cdot}$ of any labeling. Similarly, the reparametrization $\varphi_{u,v}(x_u)$, $u, v \notin A$ “outside” A does not influence it, because the optimization is performed over A only.

Considering the reparametrized potentials θ^{φ} and subtracting $\max_{x_v \in X_v} \theta_{uv}(y_u, x_v)$ in (3.3) the boundary potentials $\hat{\theta}_{uv,y_u}^{\varphi}(x_u)$ can be equivalently exchanged with

$$\begin{cases} 0, & y_u = x_u \\ \min_{x_v \in X_v} \theta_{uv}^{\varphi}(x_u, x_v) - \max_{x_v \in X_v} \theta_{uv}^{\varphi}(y_u, x_v), & y_u \neq x_u \end{cases} \quad (6.3)$$

It means that the labelings x not coinciding with y on $\partial \mathcal{V}_A$ will be “encouraged” with (typically negative) value $\Delta_{uv}^{\varphi}(x_u) := \min_{x_v \in X_v} \theta_{uv}^{\varphi}(x_u, x_v) - \max_{x_v \in X_v} \theta_{uv}^{\varphi}(y_u, x_v)$. Intuitively clear that the bigger $\Delta_{uv}^{\varphi}(x_u)$ is, the better the proposal labeling $y|_A$ comparing to $x|_A \neq y|_A$ is and hence the greater the found persistent set A^* returned by Algorithm 1 would be. We will prove correctness of this intuition formally, but first let us find the *maximal possible* value of $\Delta_{uv}^{\varphi}(x_u)$ w.r.t. the reparametrization φ , where we consider as non-zero only coordinates $\varphi_{v,u}(x_v)$, $u \in \partial \mathcal{V}_A$, $uv \in \partial \mathcal{E}_A$, $x_v \in X_v$.

Clearly

$$\begin{aligned} \Delta_{uv}^{\varphi}(x_u) &\leq \min_{x_v \in X_v} (\theta_{uv}^{\varphi}(x_u, x_v) - \theta_{uv}^{\varphi}(y_u, x_v)) \\ &= \min_{x_v \in X_v} (\theta_{uv}(x_u, x_v) + \varphi_{v,u}(x_v) - \theta_{uv}(y_u, x_v) - \varphi_{v,u}(x_v)) \\ &= \min_{x_v \in X_v} (\theta_{uv}(x_u, x_v) - \theta_{uv}(y_u, x_v)), \end{aligned} \quad (6.4)$$

hence, the right-hand-side of this inequality does not depend on the reparametrization, whereas the left-hand-side does. There is indeed such a reparametrization that turns the inequality (6.4) into equality and in this way guarantees the largest possible values of $\Delta_{uv}^{\varphi}(x_u)$ for all x_u . This, as we show below, *optimal* reparametrization is defined as

$$\varphi_{u,v}(x_v) = -\theta_{uv}(y_u, x_v), \quad (6.5)$$

which can be seen when plugging (6.5) into (6.3).

Moreover, since as we mentioned above the reparametrization “outside” an “inside” A does not influence the criterion (3.3), we can construct a single, equal for all iterations of Algorithm 1 optimal reparametrization ψ according to the rule (6.5) as

$$\psi_{u,v}(x_v) = -\theta_{uv}(y_u, x_v), \quad u \in \mathcal{V}, \quad uv \in \mathcal{E}, \quad (6.6)$$

where y is arbitrarily extended from A^0 to \mathcal{V} . Now we are ready to formulate our main result related to the reparametrization.

Let us denote $\hat{E}_{A,y}^{\varphi}$ the energy with boundary labeling defined as in Definition 3 w.r.t. the potentials θ^{φ} . Then for the reparametrization ψ defined as in (6.6) there holds

Lemma 4. *From*

$$\delta(y) \in \arg \min_{\mu \in \Lambda_A} \hat{E}_{A,y}(\mu) \quad (6.7)$$

follows $\delta(y) \in \arg \min_{\mu \in \Lambda_A} \hat{E}_{A,y}^{\psi}(\mu)$, which means: if y satisfies the persistency criterion of Corollary 1 w.r.t. potentials θ then it satisfies it w.r.t. the reparametrized potentials θ^{ψ} .

Proof. From (6.4) and (6.7) it follows that for all $uv \in \mathcal{E}_A$, $x_u \in X_u$ there holds $\hat{\theta}_{uv,y}^{\psi}(x_u) - \hat{\theta}_{uv,y}^{\psi}(y_u) \geq \hat{\theta}_{uv,y}(x_u) - \hat{\theta}_{uv,y}(y_u)$ and hence

$$\hat{E}_{A,y}^{\psi}(\mu) - \hat{E}_{A,y}^{\psi}(y) \stackrel{(6.4)}{\geq} \hat{E}_{A,y}(\mu) - \hat{E}_{A,y}(y) \geq 0 \quad (6.8)$$

for all $\mu \in \Lambda_A$. Thus $\hat{E}_{A,y}^{\psi}(y) \leq \hat{E}_{A,y}^{\psi}(\mu)$, which proves the statement of the lemma. \square

Remark 3. Lemma 4 holds for *any* polytope containing all integer solutions, i.e., $\Lambda_A \supseteq \mathcal{M}_A$ and hence it holds also when $\Lambda_A = \mathcal{M}_A$. In this case it corresponds to the non-relaxed persistency criterion provided by Theorem 1.

Let now $A_y^{\varphi,*}$ be the *largest* set containing all strongly persistent variables satisfying Corollary 1 w.r.t. the reparametrized potentials θ^{φ} and test labeling $y \in X_{\mathcal{V}}$. Let also A_y^* correspond to the trivial reparametrization $\varphi \equiv 0$.

Applying Lemma 4 to the set A_y^* leads to the following

Theorem 3. *For any test labeling $y \in X_{\mathcal{V}}$ there holds $A_y^* \subseteq A_y^{\psi,*}$.*

Proof. Same proof as in Lemma 4 applied to A_y^* . \square

Remark 4. For Potts models, where $\theta_{uv}(x_u, x_v) = \begin{cases} 0, & x_u = x_v \\ \alpha, & x_u \neq x_v \end{cases}$, the inequality (6.4) holds as equality also for trivial reparametrization $\varphi_{v,u}(x_v) = 0 \forall u, v \in \mathcal{V}, uv \in \mathcal{E}, x_v \in X_v$. For such models Algorithm 1 with trivial reparametrization delivers the same persistent set as with the optimal one (6.6).

6.2 Optimality of the Method

Theorem 2 proves optimality of Algorithm 1 w.r.t. the formulated persistency criterion provided by Theorem 1. However it does not prove optimality of the method with respect to other possible criteria and hence does not guarantee its superiority over other partial optimality techniques. There is however a recent study [29], which provides such an *optimal relaxed persistency criterion* covering all existing methods. In what follows we will introduce key notions from [29] and show that our persistency criterion coincides with the *optimal* one provided in [29] for a certain class of persistency methods, those providing only node-persistency, i.e., either eliminating all labels except one in a given node or not eliminating any.

In this section we will employ the commonly used representation of energy $E_{\mathcal{V}}(\mu)$ in a form of an inner product $\langle \theta, \mu \rangle$, where vectors of potentials θ and marginals μ belong to the vector space $\mathbb{R}^{\mathcal{I}}$ with the suitably selected dimension $\mathcal{I} = \sum_{v \in \mathcal{V}} |X_v| + \sum_{uv \in \mathcal{E}} |X_{uv}|$.

Definition 6. A mapping $p : X_{\mathcal{V}} \rightarrow X_{\mathcal{V}}$ is called (strictly) improving for the potentials θ if it is idempotent ($p(p(x)) = p(x)$) and for all $x \in X_{\mathcal{V}}$ such that $p(x) \neq x$ there holds $\langle \theta, \delta(p(x)) \rangle \leq \langle \theta, \delta(x) \rangle$ (resp. $\langle \theta, \delta(p(x)) \rangle < \langle \theta, \delta(x) \rangle$).

Following [29] we consider only *node-wise* maps of the form $p(x)_v = p_v(x_v)$, where $p_v : X_v \rightarrow X_v$ are idempotent, i.e., $p_v(p_v(x_v)) = p_v(x_v)$ for all $x_v \in X_v$. This class is already general enough to include nearly all existing techniques.

Improving mappings defines persistency due to the following proposition:

Proposition 3 (Stat.1 [29]). Let p be an improving mapping. Then there exists an optimal solution x of (2.1) such that for all $v \in \mathcal{V}$ from $p_v(i) \neq i$ follows $x_v \neq i$. In case p is strictly improving this holds for any optimal solution.

For an idempotent mapping p a linear mapping $P : \mathbb{R}^{\mathcal{I}} \rightarrow \mathbb{R}^{\mathcal{I}}$ satisfying $\delta(p(x)) = P\delta(x)$ for all $x \in X_{\mathcal{V}}$ is called its *linear extension*. A particular linear extension denoted as $[p]$ is defined as follows. For each p_v we define the matrix $P_v \in \mathbb{R}^{X_v \times X_v}$ by $P_{v,ii'} = \begin{cases} 1, & p_v(i') = i \\ 0, & p_v(i') \neq i \end{cases}$. The linear extension $P = [p]$ is given by

$$(P\mu)_v = \sum_{i' \in X_v} P_{v,ii'} \mu_v(i') = P_v \mu_v; (P\mu)_{uv} = P_u \mu_{uv} P_v^{\top}. \quad (6.9)$$

Denote by I the identity matrix. From Definition 6 follows that p is improving iff the value of

$$\begin{aligned} \min_{x \in X_{\mathcal{V}}} \langle \theta, (I - [p])\delta(x) \rangle &\equiv \min_{x \in X_{\mathcal{V}}} \langle (I - [p])^{\top} \theta, \delta(x) \rangle \\ &\equiv \min_{\mu \in \mathcal{M}_{\mathcal{V}}} \langle (I - [p])^{\top} \theta, \mu \rangle \end{aligned} \quad (6.10)$$

is zero. If additionally $p(x) = x$ for all minimizers of (6.10) then the mapping p is strictly improving.

Problem (6.10) is of the same form as energy minimization (2.1) and is therefore as hard as Problem (6.10). Its relaxation is obtained by letting μ to vary in the local polytope $\Lambda_{\mathcal{V}} \subset \mathbb{R}^{\mathcal{I}}$, an outer approximation to $\mathcal{M}_{\mathcal{V}}$.

Definition 7. Mapping $p : X_{\mathcal{V}} \rightarrow X_{\mathcal{V}}$ is $\Lambda_{\mathcal{V}}$ -improving for potentials $\theta \in \mathbb{R}^{\mathcal{I}}$ if

$$\min_{\mu \in \Lambda_{\mathcal{V}}} \langle (I - [p])^{\top} \theta, \mu \rangle = 0. \quad (6.11)$$

If additionally $[p]\mu = \mu$ for all minimizers μ of (6.11) then p is strictly $\Lambda_{\mathcal{V}}$ -improving.

Compared to (6.10), only the polytope was changed to $\Lambda_{\mathcal{V}} \supset \mathcal{M}_{\mathcal{V}}$. This implies the following simple fact:

Proposition 4. If mapping p is (strictly) $\Lambda_{\mathcal{V}}$ -improving then it is (strictly) improving.

The method presented in this work can be interpreted as considering *all-to-one* node-wise mappings p having the form

$$p_v(i) = \begin{cases} y_v, & \text{if } v \in A \\ i, & \text{if } v \notin A \end{cases} \quad (6.12)$$

for a fixed *test labeling* y . All labels in the nodes $v \in A \subset \mathcal{V}$ are mapped to y_v . Among all all-to-one (strictly) $\Lambda_{\mathcal{V}}$ -improving mappings the one with the largest set A will be called *maximal*.

Corollary 1 determines $\Lambda_{\mathcal{V}}$ -improving mappings, as stated by

Lemma 5. The relaxed persistency criterion provided by Corollary 1 with the reparametrization given by (6.6) is equivalent to Definition 7 with the improving mapping p defined as in (6.12) for a given test labeling y .

Proof. For future references we write down potentials θ^{ψ} with ψ defined by (6.6) explicitly:

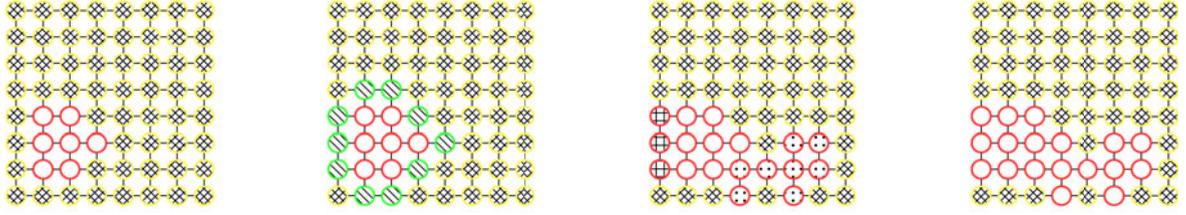
$$\begin{aligned} \theta_u^{\psi}(x_u) &= \theta_u(x_u) + \sum_{v \in \text{nb}(u)} \theta_{uv}(x_u, y_v), \theta_{uv}^{\psi}(x_u, x_v) \\ &= \theta_{uv}(x_u, x_v) - \theta_{uv}(x_u, y_v) - \theta_{uv}(y_u, x_v). \end{aligned} \quad (6.13)$$

In what follows we will show that the criteria (3.12) and (6.11) coincide. Both of them represent the local polytope relaxation of specially constructed energy minimization problems. To prove that the relaxations coincide it is sufficient to prove that the non-relaxed energies are equal.

First we write down the non-relaxed test problem (3.5) with potentials θ^{ψ} as

$$\arg \min_{x \in X_{\mathcal{V}}} \sum_{v \in \mathcal{V}} \beta_v(x_v) + \sum_{uv \in \mathcal{E}} \beta_{uv}(x_u, x_v) + \sum_{uv \in \partial \mathcal{E}_A : u \in \mathcal{V}_A} \hat{\theta}_{uv, y_u}^{\psi}(x_u) \quad (6.14)$$

with potentials β equal to θ^{ψ} on A and vanishing outside it, i.e.,



- ⊗ nodes A^t ,
- nodes $\mathcal{V} \setminus A^t$.

Boundary costs are assigned to boundary nodes ⊗ on $\partial\mathcal{V}_{A^t}$.

Having solved the inference problem:

- ⊕ nodes are fractional,
- ⊖ nodes disagree with previous labeling,
- ⊗ nodes agree.

Variables are pruned. A new set A^{t+1} is constructed.

Fig. 3. Illustration of one iteration of Algorithm 4.

$$\beta_u(x_u) = \begin{cases} \theta_u(x_u) + \sum_{v \in \text{nb}(u)} \theta_{uv}(x_u, y_v), & u \in A \\ 0, & u \in \mathcal{V} \setminus A \end{cases} \quad (6.15)$$

$$\beta_{uv}(x_u, x_v) = \begin{cases} \theta_{uv}(x_u, x_v) - \theta_{uv}(x_u, y_v) - \theta_{uv}(y_u, x_v), & u, v \in A \\ 0, & \text{otherwise.} \end{cases} \quad (6.16)$$

Border potentials $\hat{\theta}^\psi$ for $uv \in \mathcal{E}$, $u \in \mathcal{V}_A$, $v \in \mathcal{V} \setminus A$ and $x_u \neq y_u$ read:

$$\begin{aligned} \hat{\theta}_{uv, y_u}^\psi(x_u) &= \min_{x_v \in X_v} \theta_{uv}^\psi(x_u, x_v) = \\ &= \min_{x_v \in X_v} (\theta_{uv}(x_u, x_v) - \theta_{uv}(x_u, y_v) - \theta_{uv}(y_u, x_v)) \quad (6.17) \\ &= -\theta_{uv}(x_u, y_v) + \min_{x_v \in X_v} (\theta_{uv}(x_u, x_v) - \theta_{uv}(y_u, x_v)); \end{aligned}$$

for $x_u = y_u$:

$$\begin{aligned} \hat{\theta}_{uv, y_u}^\psi(y_u) &= \max_{x_v \in X_v} \theta_{uv}^\psi(y_u, x_v) = \\ &= \max_{x_v \in X_v} (\theta_{uv}(y_u, x_v) - \theta_{uv}(y_u, y_v) - \theta_{uv}(y_u, x_v)) \quad (6.18) \\ &= -\theta_{uv}(y_u, y_v). \end{aligned}$$

Note that (6.17) turns into (6.18) when $x_u = y_u$, hence it is sufficient to use only expression (6.17).

The non-relaxed version of condition (6.11) defining $\Lambda_{\mathcal{V}}$ -improving all-to-one mapping, with the labeling proposal y can be formulated as checking whether

$$y \in \arg \min_{x \in X_{\mathcal{V}}} \sum_{v \in \mathcal{V}} \gamma_v(x_v) + \sum_{uv \in \mathcal{E}} \gamma_{uv}(x_u, x_v) + \sum_{u \in \partial \mathcal{E}_A} \hat{\gamma}_{uv, y_u}(x_u) \quad (6.19)$$

with potentials γ defined as:

$$\gamma_u(x_u) = \begin{cases} \theta_u(x_u) - \theta_u(y_u), & u \in A \\ 0, & u \in \mathcal{V} \setminus A \end{cases} \quad (6.20)$$

$$\gamma_{uv}(x_u, x_v) = \begin{cases} \theta_{uv}(x_u, x_v) - \theta_{uv}(y_u, y_v), & u, v \in A \\ 0, & \text{otherwise,} \end{cases} \quad (6.21)$$

and the border term

$$\hat{\gamma}_{uv, y_u}(x_u) = \min_{x_v \in X_v} (\theta_{uv}(x_u, x_v) - \theta_{uv}(y_u, x_v)). \quad (6.22)$$

Comparing (6.20), (6.21) and (6.22) to (6.15), (6.16) and (6.17) respectively it can be seen that they can be transformed to each other by several operations, which *equally* change energies of all labelings and thus do not influence the criterions provided by Theorem 1 and [29, eq.(14)]. These operations are:

- 1) Subtract $\theta_u(y_u)$ from $\beta_u(x_u)$ for all $u \in \mathcal{V}_A$, $x_u \in X_u$.
- 2) Subtract $\theta_{uv}(y_u, y_v)$ from $\beta_{uv}(x_u, x_v)$ for all $uv \in \mathcal{E}_A$, $(x_u, x_v) \in X_u \times X_v$.
- 3) Reparametrize β with the reparametrization vector ϕ defined as

$$\phi_{u,v}(x_u) = \begin{cases} -\theta_{uv}(x_u, y_v), & u \in A \\ 0, & u \in \mathcal{V} \setminus A. \end{cases} \quad (6.23) \quad \square$$

The following theorem states that our method provably delivers the best results among the methods providing node-persistency:

Theorem 4. *Under conditions of Corollary 3, Algorithm 1 with the reparametrizations given by (6.6) finds the maximal strict $\Lambda_{\mathcal{V}}$ -improving all-to-one mapping for a given proposal labeling x^0 .*

Proof. Under condition of Corollary 3 (i.e., when on each iteration there is a unique solution μ^t) Lemma 5 guarantees equivalence of our criterion (Corollary 1 with reparametrization ψ) to Definition 7 for the strict $\Lambda_{\mathcal{V}}$ -improving all-to-one mapping. Theorem 2 states that Algorithm 1 delivers the largest set A^* satisfying this criterion, which in turn proves the theorem. \square

7 EXTENSIONS

Higher order models. Assume now we are not in the pairwise case anymore but have an energy minimization problem over a *hypergraph* $G = (\mathcal{V}, \mathcal{E})$ with $\mathcal{E} \subset \mathcal{P}(\mathcal{V})$ a set of subsets of \mathcal{V} :

$$\min_{x \in X_{\mathcal{V}}} E_{\mathcal{V}}(x) := \sum_{e \in \mathcal{E}} \theta_e(x_e). \quad (7.1)$$

All definitions, our persistency criterion and Algorithm 1 admit a straightforward generalization. Analogously to Definition 1 define for a subset of nodes $A \subset \mathcal{V}$ the

TABLE 3
Percentage of Persistent Variables Obtained by Methods [17],[20],[12],[4],[10] and Our Methods with Boundary Potentials Computed as in (3.4) (Ours Original) and as in (6.3) (Ours Optimal)

Experiment	#I	#L	#V	O	MQPBO \ [17]	Kovtun \ [20]	GRD \ [12]	Fix \ [4]	HOCR \ [10]	Ours original \ [36]	Ours optimal
teddy	1	60	168,749	2	0	†	†	†	†	0.3820	0.3820
venus	1	20	166,221	2	0	†	†	†	†	0	0
family	1	5	425,631	2	0.0432	†	†	†	†	0.0044	0.0611
pano	1	7	514,079	2	0.1247	†	†	†	†	0.2755	0.3893
Potts	12	≤12	≤424,720	2	0.1839	0.7475	†	†	†	0.9220	0.9220
side-chain	21	≤483	≤1,971	2	0.0247	†	†	†	†	0.1747	0.2558
protein-interaction	8	2	≤14,440	3	†	†	0.2603	0.2545	0.2545	0.0008	†
cell-tracking	1	2	41,134	9	†	†	†	0.1771	†	0.2966	†
geo-surf	†	†	†	†	†	†	†	†	†	0.0743	†

Notation † means inapplicability of the method. The columns #I,#L,#V,O denote the number of instances, labels, variables and the highest order of potentials respectively.

boundary nodes as

$$\partial\mathcal{V}_A := \{u \in A : \exists v \in \mathcal{V} \setminus A, \exists e \in \mathcal{E} \text{ s.t. } u, v \in e\} \quad (7.2)$$

and the boundary edges as

$$\partial\mathcal{E}_A := \{e \in \mathcal{E} : \exists u \in A, \exists v \in \mathcal{V} \setminus A \text{ s.t. } u, v \in e\}. \quad (7.3)$$

The equivalent of boundary potential in Definition 3 for $e \in \partial\mathcal{E}_A$ is

$$\hat{\theta}_{e,y}(x) := \begin{cases} \max_{\tilde{x} \in X_e : \tilde{x}|_{A \cap e} = x|_{A \cap e}} \theta_e(\tilde{x}), & x|_{A \cap e} = y|_{A \cap e} \\ \min_{\tilde{x} \in X_e : \tilde{x}|_{A \cap e} = x|_{A \cap e}} \theta_e(\tilde{x}), & x|_{A \cap e} \neq y|_{A \cap e} \end{cases} \quad (7.4)$$

Now Theorem 1, Corollary 1 and Algorithm 1 can be directly translated to the higher order case.

Tighter relaxations. Essentially, Algorithm 1 can be applied also to tighter relaxations than Λ_A , e.g., when one includes cycle inequalities [30]. One merely has to replace the local polytope Λ_A for $A \subset \mathcal{V}$ by the tighter feasible convex set:

Proposition 5. *Let the polytopes $\tilde{\Lambda}_A \supseteq \mathcal{M}_A$ satisfy $\tilde{\Lambda}_A \subseteq \Lambda_A \forall A \subseteq \mathcal{V}$. Use $\tilde{\Lambda}_{A_t}$ in place of Λ_{A_t} in Algorithm 1 and let \tilde{A}^* be the corresponding persistent set returned by the modified algorithm. Let $A_{strong}^* \subseteq A^*$ be the largest subset of strongly persistent variables identifiable by Corollary 1 subject to the relaxations $\tilde{\Lambda}_A$ and Λ_A . Then $A_{strong}^* \subseteq \tilde{A}_{strong}^*$.*

Remark 5. For approximate dual solvers for tighter relaxations like [31], [32] there are analogues of strict arc-consistency, hence these are also consistency-ascertaining solvers as in Definition 4 and we can also use these algorithms in Algorithm 1 with the obvious modifications.

Optimal reparametrization for tighter relaxations and higher order models is beyond the scope of this paper.

8 EXPERIMENTS

We tested our approach with initial and optimal reparametrizations (described in Section 6) on several datasets from

different computer vision and machine learning benchmarks, 47 problem instances overall, see Table 3. We describe each dataset and the corresponding experiments in detail below.

Competing methods. We compared our method to **MQPBO** [17], [30], **Kovtun's** method [20], Generalized Roof Duality by Kahl and Strandmark [12], Fix et al's [4] and Ishikawa's Higer Order Clique Reduction (**HOCR**) [10] algorithms. For the first two methods we used our own implementation, and for the other the freely available code of Strandmark [33]. We were unable to compare to the method of Windheuser et al. [40], because the authors do not give a description for implementing their method in the higher order case and only provide experimental evaluation for problems with pairwise potentials, where their method coincides with MQPBO [17].

Implementation details. We employed TRWS as an approximate solver for Algorithm 1 and strong tree agreement as a consistency mapping (see Proposition 2) for most of the pairwise problems. We stop TRWS once it has either arrived at (i) tree-agreement; (ii) a small duality gap of 10^{-5} ; (iii) when number of nodes with tree agreement did not increase over the last 100 iterations or (iv) overall 1,500 iterations. For the higher-order models protein-interaction, cell-tracking and geo-surf we employed CPLEX [9] as an exact linear programming solver. We have run Algorithm 1 with boundary potentials computed as in (3.3) for all problems and with boundary potentials computed with the optimal reparametrization as in (6.3) for the pairwise problems.

Datasets and Evaluation. We give a brief characterization of all datasets and report the obtained total percentage of persistent variables of our and competing methods in Table 3. The percentage of partial optimality is computed as follows: Suppose we have found a persistent labeling on set

$A \subset \mathcal{V}$. Then the percentage is $1 - \frac{\sum_{u \notin A} \log |X_u|}{\sum_{u \in \mathcal{V}} \log |X_u|}$. Note that by this formulation we take into account the size of the label

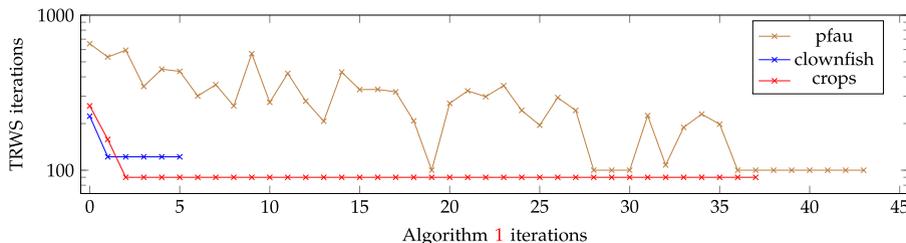


Fig. 4. Iterations needed by TRWS [18] in Algorithm 4 for three instances from the Potts dataset.

space for each node. For an uniform label space the above formula equals $\frac{|A|}{|V|}$. The latter measure was used in [35].

Remark 6. Note that in comparison to our conference paper [35], persistency results for some datasets with higher order potentials, which were solved with CPLEX are lower now. This is due to two reasons: First, we weight the size of the label space instead of simply counting the number of variables which are partially optimal. In models with nonuniform label space our method tends to find partial optimality for nodes with small label space, hence the new formula gives a smaller percentage. Second, our original research implementation contained subtle bugs which resulted in a higher number of wrongly assigned partially optimal nodes for these models. We apologize for reporting incorrect results in the experimental section of [35].

The problem instances *teddy*, *venus*, *family*, *pano*, *Potts* and *geo-surf* were made available by [13], while the datasets *side-chain* and *protein-interaction* were made available by [1].

The problem instances *teddy* and *venus* come from the **disparity estimation for stereo vision** [36]. None of the competing approaches was able to find even a single persistent variable for these datasets, presumably because of the large number of labels, whereas we labeled over one third of them as persistent in *teddy*, though none in *venus*.

Instances named *pano* and *family* come from the **photo-montage** dataset [36]. These problems have more complicated pairwise potentials than the disparity estimation problems, but less labels. For both datasets we found significantly more persistent variables than MQPBO, in particular, we were able to label more than a third of the variables in *pano*.

We also chose 12 relatively big energy minimization problems with grid structure and **Potts** interaction terms. The underlying application is a color segmentation problem previously considered in [34]. Our general approach reproduces results of [34] for the specific Potts model.

We considered also *side-chain* prediction problems in **protein folding** [41]. The datasets consist of pairwise graphical models with 32–1971 variables and 2–483 labels. The problems with fewer variables are densely connected and have very big label spaces, while the larger ones are less densely connected and have label space up to 81 variables.

The protein interaction models [11] aim to find the subset of proteins, which interact with each other. Roof-duality based methods, i.e., Fix et al., GRD, HOCR [4], [10], [12] gave around a quarter of persistent labels. This is the only dataset where our methods gives worse results. Note that for higher-order models we do not provide an optimal reparametrization and hence our method is not provably better than the competitors. We consider this as a direction for future work.

The **cell tracking** problem consists of a binary higher order graphical model [16]. Given a sequence of microscopy images of a growing organism, the aim is to find the lineage tree of all cells. For implementation reasons we were not able to solve *cell-tracking* dataset with Ishikawa's [10] method. However Fix [4] reports that his method outperforms Ishikawa's method [10]. Other methods are not applicable even theoretically.

Last, we took the higher order multi-label **geometric surface labeling problems** (denoted as *geo-surf* in Table 3) from [8]. The only instance having an integrality gap has 968 variables with seven labels each and has ternary terms. Note that MQPBO cannot handle ternary terms, Fix et al.'s [4] Ishikawa's [10] methods and the

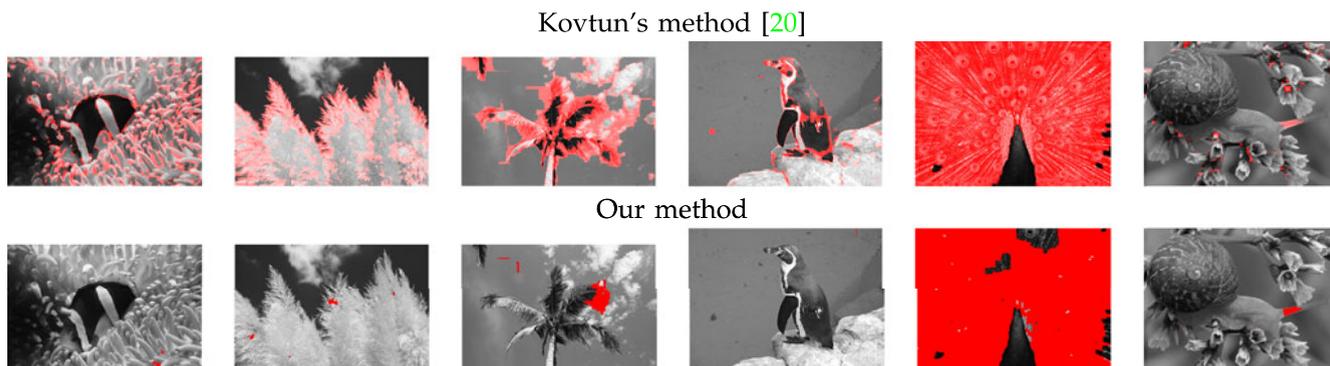


Fig. 5. Comparison between Kovtun's Method [20] and Our Method. The red area denotes pixels which could not be labelled persistently. Contrary to ours the Kovtun's method allows to eliminate separate labels, which is denoted by different intensity of the red color: the more intensive is red, the less labels were eliminated.

generalized roof duality method by Strandmark and Kahl [12] cannot handle more than two labels. Hence we report our results without comparison.

Runtime. The runtime of our algorithm mainly depends on the speed of the underlying solver for the local polytope relaxation. Currently there seems to be no general rule regarding the runtime of our algorithm, neither in the number of Algorithm 1-iterations nor in the number of TRWS [18]-iterations. We show three iteration counts for instances of the Potts dataset in Fig. 4.

Exemplary pictures comparing the pixels optimally labelled between Kovtun's method [20] and our method for some Potts-models can be seen in Fig. 5.

9 CONCLUSION AND OUTLOOK

We have presented a novel method for finding persistent variables for undirected graphical models. Empirically it outperforms all tested approaches with respect to the number of persistent variables found on every single dataset. Our method is general: it can be applied to graphical models of arbitrary order and type of potentials. Moreover, there is no fixed choice of convex relaxation for the energy minimization problem and also approximate solvers for these relaxations can be employed in our approach.

In the future we plan to significantly speed-up the implementation of our method and consider finer persistency criteria, as done in [29], where the subset-to-one class of persistency conditions was introduced, but no efficient algorithm for finding persistency in this class was proposed.

ACKNOWLEDGMENTS

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.

REFERENCES

- [1] (2011). The probabilistic inference challenge (PIC2011) [Online]. Available: <http://www.cs.huji.ac.il/project/PASCAL/>
- [2] E. Boros and P. L. Hammer, "Pseudo-Boolean optimization," *Discr. Appl. Math.*, vol. 123, nos. 1–3, pp. 155–225, 2002.
- [3] J. Desmet, M. D. Maeyer, B. Hazes, and I. Lasters, "The dead-end elimination theorem and its use in protein side-chain positioning," *Nature*, vol. 356, no. 6369, pp. 539–542, Apr. 1992.
- [4] A. Fix, A. Gruber, E. Boros, and R. Zabih, "A graph cut algorithm for higher-order Markov random fields," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2011, pp. 1020–1027.
- [5] A. Globerson and T. Jaakkola, "Fixing max-product: Convergent message passing algorithms for MAP LP-relaxations," in *Proc. Adv. Neural Inf. Process. Syst.*, 2008, pp. 553–560.
- [6] P. Hammer, P. Hansen, and B. Simeone, "Roof duality, complementation and persistency in quadratic 0-1 optimization," *Math. Program.*, vol. 28, pp. 121–155, 1984.
- [7] T. Hazan and A. Shashua, "Norm-product belief propagation: Primal-dual message-passing for approximate inference," *IEEE Trans. Inf. Theory*, vol. 56, no. 12, pp. 6294–6316, Dec. 2010.
- [8] D. Hoiem, A. A. Efros, and M. Hebert, "Recovering surface layout from an image," *Int. J. Comput. Vis.*, vol. 75, no. 1, pp. 151–172, Oct. 2007.
- [9] (2014). ILOG, Inc. ILOG CPLEX: High-performance software for mathematical programming and optimization [Online]. Available: <http://www.ilog.com/products/cplex/>
- [10] H. Ishikawa, "Transformation of general binary MRF minimization to the first-order case," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 6, pp. 1234–1249, Jun. 2011.
- [11] A. Jaimovich, G. Elidan, H. Margalit, and N. Friedman, "Towards an integrated protein-protein interaction network: A relational Markov network approach," *J. Comput. Biol.*, vol. 13, no. 2, pp. 145–164, 2006.
- [12] F. Kahl and P. Strandmark, "Generalized roof duality," *Discr. Appl. Math.*, vol. 160, no. 16–17, pp. 2419–2434, 2012.
- [13] 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 problem," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2013, pp. 1328–1335.
- [14] J. H. Kappes, B. Savchynskyy, and C. Schnörr, "A bundle approach to efficient MAP-inference by Lagrangian relaxation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2012, pp. 1688–1695.
- [15] J. H. Kappes, M. Speth, G. Reinelt, and C. Schnörr, "Towards efficient and exact MAP-inference for large scale discrete computer vision problems via combinatorial optimization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2013, pp. 1752–1758.
- [16] B. X. Kausler, M. Schiegg, B. Andres, M. S. Lindner, U. Köthe, H. Leitte, J. Wittbrodt, L. Hufnagel, and F. A. Hamprecht, "A discrete chain graph model for 3d+t cell tracking with high misdetection robustness," in *Proc. 12th Eur. Conf. Comput. Vis.*, 2012, pp. 144–157.
- [17] P. Kohli, A. Shekhovtsov, C. Rother, V. Kolmogorov, and P. Torr, "On partial optimality in multi-label MRFs," in *Proc. 25th Int. Conf. Mach. Learn.*, New York, NY, USA, 2008, pp. 480–487.
- [18] V. Kolmogorov, "Convergent tree-reweighted message passing for energy minimization," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 10, pp. 1568–1583, Oct. 2006.
- [19] N. Komodakis, N. Paragios, and G. Tziritas, "MRF energy minimization and beyond via dual decomposition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 3, pp. 531–552, Mar. 2011.
- [20] I. Kovtun, "Partial optimal labeling search for a NP-hard subclass of (max,+) problems," vol. 2781, pp. 402–409, 2003.
- [21] I. Kovtun, "Sufficient condition for partial optimality for (max, +) labeling problems and its usage," *Control Syst. Comput.*, vol. 2, pp. 35–42, 2011.
- [22] G. L. Nemhauser and L. E. Trotter, "Vertex packings: Structural properties and algorithms," *Math. Program.*, vol. 8, pp. 232–248, 1975.
- [23] C. Rother, V. Kolmogorov, V. S. Lempitsky, and M. Szummer, "Optimizing binary MRFs via extended roof duality," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2007, pp. 1–8.
- [24] B. Savchynskyy, J. H. Kappes, S. Schmidt, and C. Schnörr, "A study of Nesterov's scheme for Lagrangian decomposition and MAP labeling," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2011, pp. 1817–1823.
- [25] B. Savchynskyy, J. H. Kappes, P. Swoboda, and C. Schnörr, "Global MAP-optimality by shrinking the combinatorial search area with convex relaxation," in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 1950–1958.
- [26] B. Savchynskyy, S. Schmidt, J. H. Kappes, and C. Schnörr, "Efficient MRF energy minimization via adaptive diminishing smoothing," in *Proc. Uncertainty Artif. Intell.*, 2012, pp. 746–755.
- [27] M. Schlesinger, "Syntactic analysis of two-dimensional visual signals in the presence of noise," *Kibernetika*, vol. 12, no. 4, pp. 113–130, 1976.
- [28] S. Schmidt, B. Savchynskyy, J. H. Kappes, and C. Schnörr, "Evaluation of a first-order primal-dual algorithm for MRF energy minimization," in *Proc. 8th Int. Conf. Energy Minimization Methods Comput. Vis. Pattern Recog.*, 2011, pp. 89–103.
- [29] A. Shekhovtsov, "Maximum persistency in energy minimization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2014, pp. 1162–1169.
- [30] D. Sontag, "Approximate inference in graphical models using LP relaxations," Ph.D. dissertation, Dept. Elect. Eng. Comput. Sci., Massachusetts Inst. Technol., Cambridge, MA, USA, 2010.
- [31] D. Sontag, D. K. Choe, and Y. Li, "Efficiently searching for frustrated cycles in MAP inference," in *Proc. 28th Conf. Annu. Conf. Uncertainty Artif. Intell.*, 2012, pp. 795–804.
- [32] D. Sontag, T. Meltzer, A. Globerson, T. Jaakkola, and Y. Weiss, "Tightening lp relaxations for map using message passing," in *Proc. 24th Annu. Conf. Uncertainty Artif. Intell.*, 2008, pp. 503–510.
- [33] P. Strandmark, Generalized roof duality [Online]. Available: <http://www.maths.lth.se/matematiklth/personal/petter/pseudoboolen.php>

- [34] P. Swoboda, B. Savchynskyy, J. H. Kappes, and J. Schnörr, "Partial optimality via iterative pruning for the Potts model," in *Proc. 4th Int. Conf. Scale Space Variational Methods Comput. Vis.*, pp. 477–488, 2013.
- [35] P. Swoboda, B. Savchynskyy, J. H. Kappes, and C. Schnörr, "Partial optimality by pruning for MAP-inference with general graphical models," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2014, pp. 1170–1177.
- [36] R. Szeliski, R. Zabih, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. F. Tappen, and C. Rother, "A comparative study of energy minimization methods for Markov random fields with smoothness-based priors," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 6, pp. 1068–1080, Jun. 2008.
- [37] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky, "MAP estimation via agreement on trees: Message-passing and linear programming," *IEEE Trans. Inf. Theor.*, vol. 51, no. 11, pp. 3697–3717, Nov. 2005.
- [38] M. J. Wainwright and M. I. Jordan, "Graphical models, exponential families, and variational inference," *Found. Trends Mach. Learn.*, vol. 1, nos. 1/2, pp. 1–305, Jan. 2008.
- [39] T. Werner, "A linear programming approach to max-sum problem: A review," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 7, pp. 1165–1179, Jul. 2007.
- [40] T. Windheuser, H. Ishikawa, and D. Cremers, "Generalized roof duality for multi-label optimization: Optimal lower bounds and persistency," in *Proc. 12th Eur. Conf. Comput. Vis.*, Firenze, Italy, Oct. 2012, pp. 400–413.
- [41] C. Yanover, O. Schueler-Furman, and Y. Weiss, "Minimizing and learning energy functions for side-chain prediction," *J. Comput. Biol.*, vol. 15, no. 7, pp. 899–911, 2008.



Paul Swoboda graduated in mathematics from Heidelberg University, Germany, in 2011. Since then, he has been a predoctoral fellow at Heidelberg University. His main interests include variational problems in image analysis and combinatorial and variational methods for solving them.



Alexander Shekhovtsov received the MS degree in applied mathematics from the National Technical University of Ukraine "Kiev Polytechnic Institute." In 2013, he defended his PhD thesis on exact and partial energy minimization in computer vision. From 2001 to 2005, he was under the supervision of M.I. Schlesinger. From 2005 to 2013, he joined the Center for Machine Perception, Czech Technical University in Prague. Since Feb. 2014, he has been a postdoc at the Institute for Computer Graphics and Vision, Graz University of Technology.



Jörg Hendrik Kappes received the diploma degree in computer science from the University of Mannheim in 2005 and the PhD degree from Heidelberg University in 2011. Since 2011, he has been a postdoc at Heidelberg University. He is one of the major developers of OpenGM2, a free library for inference in discrete graphical models. His research interests include discrete optimization, computer vision, and probabilistic graphical models.



Christoph Schnörr received the degrees from the Technical University of Karlsruhe, Germany. After moving to the Universities of Hamburg and Mannheim, he joined Heidelberg University where he is currently heading the Image and Pattern Analysis Group, co-directing the Heidelberg Collaboratory for Image Processing, and coordinating the Research Training Group on Probabilistic Graphical Models and Applications in Image Analysis funded by the German Research Foundation (DFG). His research interests include mathematical models and algorithm design for image and pattern analysis.



Bogdan Savchynskyy graduated from the National Technical University of Ukraine "Kiev Polytechnic Institute" in 2002. In 2007, he defended the PhD thesis at the National Academy of Sciences of Ukraine. Since 2009, has been a postdoc at Heidelberg University, Germany. His main interests include optimization problems for image and video analysis, in particular inference and learning with image grammars, and Markov random fields.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.