

Efficient Semidefinite Branch-and-Cut for MAP-MRF Inference

Peng Wang · Chunhua Shen · Anton van den Hengel · Philip H. S. Torr

the date of receipt and acceptance should be inserted later

Abstract We propose a Branch-and-Cut (B&C) method for solving general MAP-MRF inference problems. The core of our method is a very efficient bounding procedure, which combines scalable semidefinite programming (SDP) and a cutting-plane method for seeking violated constraints. In order to further speed up the computation, several strategies have been exploited, including model reduction, warm start and removal of inactive constraints.

We analyze the performance of the proposed method under different settings, and demonstrate that our method either outperforms or performs on par with state-of-the-art approaches. Especially when the connectivities are dense or when the relative magnitudes of the unary costs are low, we achieve the best reported results. Experiments show that the proposed algorithm achieves better approximation than the state-of-the-art methods within a variety of time budgets on challenging non-submodular MAP-MRF inference problems.

1 Introduction

Markov Random Fields (MRFs) have been used to model a variety of problems in computer vision, including semantic image segmentation, restoration, 3D reconstruction and

stereo matching, amongst a lot of others. Finding the maximum a posteriori (MAP) solution to *general* MRF problems is typically NP-hard, however, and many approaches have been proposed to solve such problems, approximately or exactly (see [1, 2, 3, 4, 5] for comparative studies).

Graph cuts based methods [6, 7, 8] have been applied to a number of MAP problems in computer vision. Binary MRFs with submodular pairwise potentials can be solved exactly and efficiently by graph cuts. The QPBO [8, 9] algorithm can obtain a part of the globally optimal solution for some non-submodular, binary, pairwise MRFs, but its performance degrades as the portion of non-submodular potentials increases and in the case of highly connected graphs or weak unary potentials [9]. For multi-label MRFs, expansion move and swap move algorithms [6] have been employed to find a strong local optimum with the property that no expansion move (swap move) can decrease the energy. However, this optimality is guaranteed only if the binary subproblem at each iteration is solved globally. In particular, the binary subproblem is typically required to be submodular to guarantee this optimality.

Another class of popular inference approaches is based on message passing. Max-product belief propagation [10, 11, 12, 13] obtains exact solutions to tree-structured graphs, single-cycle graphs [14, 15], and maximum weight matching on bipartite graphs [16]. For graphs with cycles, approximate solutions can also be obtained using max-product belief propagation, but there is no convergence guarantee. Tree-reweighted max-product (TRW) message-passing [17, 18, 19] differs from belief propagation in that it maintains a lower-bound to the minimum energy and can be used to measure the quality of approximate solutions. TRW message passing is proved [19] to be exact for binary submodular MRFs. Kolmogorov [18] proposed an improved version of TRW, called TRW-S, with convergence guarantee that TRW does not have.

P. Wang
University of Adelaide, SA 5005, Australia.

C. Shen (✉)
University of Adelaide, SA 5005, Australia;
and Australian Centre for Robotic Vision.
E-mail: chunhua.shen@adelaide.edu.au

A. van den Hengel
University of Adelaide, SA 5005, Australia;
and Australian Centre for Robotic Vision.

P. H. S. Torr
University of Oxford, United Kingdom.

Many aforementioned approaches have underlying connections to the standard linear programming (LP) relaxation optimizing over *local marginal polytope* [20, 21, 22]. The ordinary max-product updates [10] can be considered as a Lagrangian approach to the dual of the standard LP relaxation, for tree-structured graphs. However, this relationship cannot be generalized to graphs with cycles. For binary pairwise graphs, QPBO and TRW message passing exactly solve the standard LP relaxation [19]. The standard LP relaxation can achieve exact solutions for tree-structured graphs or binary graphs with submodular potentials. It is proved in [3] that the standard LP relaxation dominates second-order cone programming (SOCP) relaxation [23] and quadratic programming (QP) relaxation [24].

Standard LP algorithms such as interior-point methods are usually computationally inefficient for large-scale MAP problems, as they do not exploit the graph structure. To this end, a number of specific approaches to solve LP relaxation for MAP problems have been proposed, such as block coordinate descent methods [18, 25], subgradient descent methods [26, 27, 28], bundle methods [29], proximal methods [30], ADMM [31, 32, 33] and smoothing methods [30, 34, 35, 36].

However, it has been shown in [37, 38, 39, 40] that the standard LP relaxation is still not tight enough for many hard MAP problems in real applications. Especially, the above-mentioned methods usually perform poorly for the class of densely-connected graphs with weak unary potentials and a large portion of non-submodular pairwise potentials (*which are of major interests in this paper*). Because there are a large number of (potentially frustrated) long cycles in densely-connected graphs, the standard LP relaxation is likely to provide loose bounds. Two approaches can be adopted to alleviate this issue. The first direction is based on LP relaxation with high-order consistency constraints [37, 38, 39, 40, 41, 42]; and the second one is to use semidefinite programming (SDP) relaxation [43, 44, 45, 46, 47, 48, 49, 50, 51].

As the standard LP relaxation only enforces edge consistency constraints over pseudo-marginals of variables, it can be loose due to the existence of violated high-order consistency constraints. Cluster based methods [37, 39, 40, 42] can be used to tighten the standard LP relaxation by incrementally adding high-order consistency constraints. Methods for searching violated high-order constraints include dictionary enumeration for triplets or other short cycles [39, 40] and separation algorithms for long cycles [42, 52].

SDP relaxation provides an alternative tighter bound for MAP-MRF inference problems, compared with LP relaxation. SDP is of great importance in developing approximation algorithms for some NP-hard optimization problems, and usually provides more accurate solutions. In particular for the classical maxcut problem, it achieves the best approximation ratio of 0.879 [53]. Primal-dual interior-point

methods [54, 55, 56, 57, 58] are considered as state-of-the-art general SDP solvers, which in the worst case require $\mathcal{O}(m^3 + mn^3h^3 + m^2n^2h^2)$ arithmetic operations and $\mathcal{O}(m^2 + n^2h^2)$ memory requirement to solve SDP relaxation to an MAP problem with n nodes, h states per node and m linear constraints. The exponent in the polynomial complexity bound is so great as to preclude practical applications of interior-point methods to even medium sized problems. *This has significantly hampered the use of SDP relaxations in MAP-MRF inference.*

Some approaches [47, 48, 51, 59] solve SDP relaxation efficiently based on the low-rank approximation of p.s.d. matrix: $\mathbf{X} = \mathbf{Y}\mathbf{Y}^\top$ (such that p.s.d. constraints are eliminated). However, these methods need to solve a sequence of non-convex problems, and may get stuck in a local-optimal point. Multivariate weight updates based methods [60, 61, 62, 63, 64, 65] solve specific SDP problems inexactly, but need many iterations to converge to accurate solutions. Augmented Lagrangian methods [50, 66, 67, 68] have also been proposed for solving SDP problems, which can be viewed as an instance of gradient descent methods [69] and may have a slow convergence rate.

It is shown in [22] that MAP LP relaxation with local consistency constraints and conventional SDP relaxation without considering local consistency constraints are mutually incomparable, which means that neither of them is tighter than the other. Thus combining conventional SDP relaxation and standard/high-order local consistency constraints together would produce an even tighter bound. Unfortunately, this combination results in a very challenging optimization problem, not only because of the computational inefficiency of SDP itself, but also due to the large number of linear constraints arising from standard/high-order LP relaxation.

In this paper, we propose an efficient SDP approximation/bounding approach, which is faster than state-of-the-art competing methods and able to handle a large number of linear constraints. A Branch-and-Cut (B&C) method is developed based on this SDP bounding procedure. Our approach can be used to achieve either the exact solution or an accurate approximation to general MAP-MRF inference problems with a time budget, and does so at a lower computational cost than state-of-the-art competing methods. The main contributions of this paper are as follows.

- We present an approximate inference approach based on a scalable SDP algorithm [49] and cutting-plane. The proposed formulation minimizes the energy over the intersection of the semidefinite and polyhedral outer-bounds arising from SDP and standard/high-order LP relaxation respectively. Such optimization schemes result in SDP problems with a large number of linear constraints m . Significantly, our proposed SDP method scales *linearly* in m (i.e., $\mathcal{O}(m)$)—a sharp contrast to standard

interior-point methods with complexity of $\mathcal{O}(m^3)$. The method proceeds by incrementally adding violated constraints until the required solution quality is achieved. Our method is much faster than interior-point methods while still maintains a comparable lower-bound of the minimum energy.

- The proposed SDP approximate/bounding method is embedded into a B&C framework for exactly solving MAP-MRF inference problems. We also introduce a few techniques to optimize the bounding and branching procedures, including model reduction, warm start and removal of inactive constraints.
- We analyze the performance on a variety of MRF graphic models, which demonstrates that our method performs better than recent competing approaches, especially when the graph connectivity is dense and/or the unary potentials are weak.

Related work We review some research work that is closely related to ours. The proposed approach is motivated by the line of our prior work in [49, 70]. While the work in [49, 70] focuses on distance metric learning and image segmentation, our work here considers approximate and exact methods for general MAP-MRF inference problems.

Different B&B algorithms have been proposed for MAP-MRF inference. Sun *et al.* [71] proposed a B&B method (refer to as MPLP-BB) based on LP relaxation. DAOOPT [72] is another B&B method based on LP relaxation, which combines several sophisticated techniques, including And/Or search spaces, the mini-bucket approach and stochastic local search. In our experiments, MPLP-BB and DAOOPT are outperformed by our method. Peng *et al.* [48] developed an approximate inference method for optimizing over the intersection of semidefinite bound and local marginal polytope. Note that the optimization technique in their work is different from ours. There, the p.s.d. constraint and local marginal polytope are separated using dual decomposition, and then the SDP subproblem is solved using non-convex QP, which results in a local optimum. On the contrary, our method processes the p.s.d. and linear constraints altogether and optimizes a convex problem with a guarantee of being globally optimal.

SDP-based B&B or B&C approaches have also been widely studied for integer/continuous quadratic problems. Krislock *et al.* [73] introduced an SDP-based branch-and-bound (B&B) method to solve Max-Cut problems. An experimental comparison was performed in [74] between LP- and SDP-based B&B methods for graph bisection problems, which showed that SDP-based methods are better for graphs with a few thousands nodes. The algorithm in [75] solves *real-valued* nonconvex quadratic programs using a finite B&B algorithm in which the bound of each branch is computed by semidefinite relaxation. The work in [76] introduced an exact solver for binary quadratic programs, which combines

DC (difference of convex functions) algorithms and SDP-based B&B approaches. Mars and Schewe [77] proposed an SDP-based B&B method for mixed-integer semidefinite programs. SDP-based cutting plane algorithms were studied in [78, 79, 80, 81, 82] for binary quadratic programs, in which linear constraints arising from cut polytope were used to strengthen the SDP relaxation. In contrast to the above approaches, our algorithm aims to solve the MAP inference problems for discrete graphical models, and uses a number of specific speedup strategies.

Notation Table 1 lists the notation used in this paper:

\mathbf{X}	A matrix (bold upper-case letters).
\mathbf{x}	A column vector (bold lower-case letters).
\mathbb{R}	The set of real numbers.
\mathcal{S}^n	The set of $n \times n$ symmetric matrices.
\mathcal{S}_+^n	The cone of positive semidefinite (p.s.d.) $n \times n$ matrices.
$\mathbf{X} \succeq 0$	The matrix \mathbf{X} is positive semidefinite.
\leq, \geq	Inequality between scalars or element-wise inequality between column vectors.
$\mathbf{1}(\cdot)$	Indicator function, 1 if the statement is true and 0 otherwise.
$\text{trace}(\cdot)$	Trace of a matrix.
$\text{rank}(\cdot)$	Rank of a matrix.
$\text{diag}(\mathbf{X})$	The main diagonal vector of the matrix \mathbf{X} .
$\text{diag}(\mathbf{x})$	A diagonal matrix consisting of the input vector \mathbf{x} as diagonal elements.
$\ \cdot\ _1, \ \cdot\ _2$	ℓ_1 and ℓ_2 norm of a vector.
$\ \cdot\ _F$	Frobenius-norm of a matrix.
$\langle \cdot, \cdot \rangle$	Inner product of two matrices.
$\nabla f(\cdot)$	The first-order derivatives of function $f(\cdot)$.
$\lfloor x \rfloor$	The nearest integer less than or equal to x .

Table 1: Notation.

The remaining of this paper is organized as follows. We first introduce the basic concepts of LP and SDP relaxation to MAP problems in Section 2. Then our main contributions, the overall branch-and-cut method for energy minimization and the SDP bounding approach with cutting-plane are discussed in Sections 3 and 4 respectively. The performance of our algorithms is evaluated experimentally in Section 5.

2 LP and SDP Relaxation to MAP-MRF Inference Problems

Suppose that an MAP-MRF inference model is characterized by an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} represents the set of n nodes (which is reused to represent the set of node indexes $\{1, 2, \dots, n\}$) and \mathcal{E} represents the set of edges. In this paper, we only consider unary and pairwise potentials, and so MAP-MRF inference problems can be expressed as the following energy minimization problem:

$$\min_{\mathbf{x} \in \mathbb{Z}^n} E(\mathbf{x}) := \sum_{p \in \mathcal{V}} \theta_p(x_p) + \sum_{p < q, (p,q) \in \mathcal{E}} \theta_{p,q}(x_p, x_q), \quad (1)$$

where $\theta_p(x_p)$ and $\theta_{p,q}(x_p, x_q)$ represent the unary and pairwise potentials respectively. $\mathcal{Z} = \{1, 2, \dots, h\}$ refers to the set of h possible states. Without loss of generality, we assume that all nodes have the same number of possible states.

It is well known that the energy minimization problem (1) is equivalent to the following LP problem:

$$\min_{\{\mathbf{y}, \mathbf{Y}\} \in \mathcal{M}(\mathcal{G}, \mathcal{Z})} E(\mathbf{y}, \mathbf{Y}) := \sum_{p \in \mathcal{V}, i \in \mathcal{Z}} \theta_p(i) y_{p,i} + \sum_{p < q, (p,q) \in \mathcal{E}, i, j \in \mathcal{Z}} \theta_{p,q}(i, j) Y_{pi,qj} \quad (2)$$

where the set $\mathcal{M}(\mathcal{G}, \mathcal{Z})$ refers to the *marginal polytope* with respect to graph \mathcal{G} and the state space \mathcal{Z} :

$$\mathcal{M}(\mathcal{G}, \mathcal{Z}) := \left\{ \mathbf{y}, \mathbf{Y} \mid \begin{array}{l} \exists \Pr(\mathbf{x}), \text{ s.t.} \\ y_{p,i} = \Pr(x_p = i), \quad \forall p \in \mathcal{V}, i \in \mathcal{Z}, \\ Y_{pi,qj} = \Pr(x_p = i, x_q = j), \quad \forall (p, q) \in \mathcal{E}, i, j \in \mathcal{Z}. \end{array} \right\} \quad (3)$$

In the above definition, $\{y_{p,i}\}_{i \in \mathcal{Z}}$ and $\{Y_{pi,qj}\}_{i,j \in \mathcal{Z}}$ refer to singleton and pairwise pseudomarginals with respect to each node $p \in \mathcal{V}$ and each edge $(p, q) \in \mathcal{E}$ respectively, which are globally realizable by some joint distribution $\Pr(\cdot)$ over $\mathbf{x} \in \mathcal{Z}^n$. Each vertex (extreme point) of $\mathcal{M}(\mathcal{G}, \mathcal{Z})$ is integer ($\{0, 1\}$ -valued) and corresponds to a valid assignment $\mathbf{x} \in \mathcal{Z}^n$. So there is always a binary optimal solution to the LP problem (2).

Although the above LP problem only contains $\mathcal{O}(|\mathcal{E}| + |\mathcal{V}|)$ number of variables, it generally requires exponential number of linear constraints to describe $\mathcal{M}(\mathcal{G}, \mathcal{Z})$. To address this difficulty, the standard LP relaxation [20, 21, 22] optimizes the objective function of (2) over a tractable set, called *local marginal polytope*:

$$\mathcal{M}_L(\mathcal{G}, \mathcal{Z}) := \left\{ \mathbf{y}, \mathbf{Y} \mid \begin{array}{l} \mathbf{y} \geq \mathbf{0}; \mathbf{Y} \geq \mathbf{0}; \\ \sum_{i \in \mathcal{Z}} y_{p,i} = 1, \quad \forall p \in \mathcal{V}; \\ \sum_{j \in \mathcal{Z}} Y_{pi,qj} = y_{p,i}, \quad \forall (p, q) \in \mathcal{E}, i \in \mathcal{Z}. \end{array} \right\}, \quad (4)$$

which is a convex polyhedral outer-bound on $\mathcal{M}(\mathcal{G}, \mathcal{Z})$. That is, $\mathcal{M}(\mathcal{G}, \mathcal{Z}) \subseteq \mathcal{M}_L(\mathcal{G}, \mathcal{Z})$. Its number of linear constraints grows linearly with respect to the graph size but quadratically with the size of \mathcal{Z} . It is known (see [22] for example) that any vertex of $\mathcal{M}(\mathcal{G}, \mathcal{Z})$ is also a vertex of $\mathcal{M}_L(\mathcal{G}, \mathcal{Z})$, and for a graph \mathcal{G} with cycles, $\mathcal{M}_L(\mathcal{G}, \mathcal{Z})$ also contains fractional vertexes lying outside $\mathcal{M}(\mathcal{G}, \mathcal{Z})$. Consequently, we have the following important property of the standard LP relaxation (see [22]):

Theorem 1 (Optimality guarantee of LP relaxation)

If the solution to the LP relaxation: $\min_{\{\mathbf{y}, \mathbf{Y}\} \in \mathcal{M}_L(\mathcal{G}, \mathcal{Z})} E(\mathbf{y}, \mathbf{Y})$ is integer, it is the exact MAP solution.

The edge consistency is preserved in $\mathcal{M}_L(\mathcal{G}, \mathcal{Z})$, but high-order consistency is not enforced. Cluster based methods [39, 40, 42] iteratively add to LP relaxation high-order consistency

constraints, and provide incrementally improved lower-bounds to the minimal energy: $\min_{\mathbf{x} \in \mathcal{Z}^n} E(\mathbf{x})$.

Alternatively, SDP relaxation provides a different convex outer-bound on $\mathcal{M}(\mathcal{G}, \mathcal{Z})$. Consider a symmetric matrix consisting of \mathbf{y} and \mathbf{Y} :

$$\Omega(\mathbf{y}, \mathbf{Y}) := \begin{bmatrix} 1 & y_{1,1} & \cdots & y_{p,i} & \cdots & y_{q,j} & \cdots & y_{n,h} \\ y_{1,1} & Y_{11,11} & \cdots & Y_{11,pi} & \cdots & Y_{11,qj} & \cdots & Y_{11,nh} \\ \vdots & \vdots & & \vdots & & \vdots & & \vdots \\ y_{pi} & Y_{pi,11} & \cdots & Y_{pi,pi} & \cdots & Y_{pi,qj} & \cdots & Y_{pi,nh} \\ \vdots & \vdots & & \vdots & & \vdots & & \vdots \\ y_{qj} & Y_{qj,11} & \cdots & Y_{qj,pi} & \cdots & Y_{qj,qj} & \cdots & Y_{qj,nh} \\ \vdots & \vdots & & \vdots & & \vdots & & \vdots \\ y_{nh} & Y_{nh,11} & \cdots & Y_{nh,pi} & \cdots & Y_{nh,qj} & \cdots & Y_{nh,nh} \end{bmatrix}, \quad (5)$$

where $Y_{pi,qj} := y_{p,i} \cdot y_{q,j}$, $\forall (p, q) \notin \mathcal{E}, i, j \in \mathcal{Z}$. At any vertex of $\mathcal{M}(\mathcal{G}, \mathcal{Z})$, $\Omega(\mathbf{y}, \mathbf{Y}) = \Omega(\mathbf{y}, \mathbf{y}\mathbf{y}^\top) = \begin{bmatrix} 1 & \mathbf{y}^\top \\ \mathbf{y} & \mathbf{y}\mathbf{y}^\top \end{bmatrix} \succeq \mathbf{0}$ is a rank-1 positive semidefinite matrix.

By discarding the non-convex rank-1 constraint, we obtain the following convex outer-bound on $\mathcal{M}(\mathcal{G}, \mathcal{Z})$:

$$\mathcal{P}(\mathcal{G}) := \{\mathbf{y}, \mathbf{Y} \mid \Omega(\mathbf{y}, \mathbf{Y}) \succeq \mathbf{0}\}, \quad (6)$$

which generally contains curved boundaries. It is shown in [22] that local marginal polytope $\mathcal{M}_L(\mathcal{G}, \mathcal{Z})$ and the semidefinite outer-bound $\mathcal{P}(\mathcal{G})$ are mutually incomparable, which means that neither of them dominates the other. We attempt to combine the local marginal polytope and semidefinite outer-bound in this work such that a tighter bound can be achieved. To solve the resulted SDP problem, which may have a large number of constraints, scalable SDP solvers are developed.

In general, both LP and SDP relaxation do not provide feasible integer solutions to the original non-submodular energy minimization problem. Thus a rounding procedure is needed to transform a fractional solution given by LP/SDP relaxation to a feasible integer solution. A simple rounding procedure for both LP/SDP relaxation is to round each variable independently based on the singleton pseudo-marginals \mathbf{y} :

$$x_p = \arg \max_{i \in \mathcal{Z}} y_{p,i}, \quad p = 1, 2, \dots, n. \quad (7)$$

Better rounding methods would consider the labelling over all variables jointly, which can be done by considering the reparameterized objective [18] or pairwise pseudo-marginals \mathbf{Y} conditioned by already labelled variables. For max-cut problems, the random hyperplane rounding approach [53] for SDP relaxation delivers a bounded expected objective value. In this paper, the simple rounding procedure (7) is adopted for our SDP solver, followed by a post-processing procedure based on ICM [83] (see Section 4.5).

3 Branch-and-Cut for Energy Minimization

In this section, the overall algorithm of B&C for the energy minimization problem (1) is given in Algorithm 1 and then

the subproblem selection and branching strategies are presented.

3.1 The Branch-and-Bound Method

First, we briefly revisit B&B methods [84, 85]. As a method for global optimization of nonconvex problems, B&B methods rely on two procedures which compute upper- and lower-bounds on the global optimum. Considering a minimization problem, the upper-bound can be obtained by choosing any point in the feasible set. A local search method is usually performed to improve the upper-bound. On the other hand, the lower-bound can be computed from a convex relaxation to the original non-convex problem.

At each iteration of B&B (Algorithm 1), a subproblem is selected from the priority queue \mathcal{Q} (Step 1) and its upper-/lower-bounds are computed (Step 2). In Step 3, the global upper-bound (gub) is updated in each iteration as the minimum of the upper-bounds over all branches, while the global lower-bound (glb) is updated as the minimum of the lower-bounds over all leaf branches. All subproblems with a lower-bound not smaller than the global upper-bound are pruned in Step 4. If the selected subproblem cannot be pruned, its feasible set \mathcal{D} is separated into at least two convex sets (Step 5). B&B algorithms terminate when the global upper-bound and lower-bound are sufficiently close to each other.

There are several essential components in B&B: bounding ($\text{bound}(\mathcal{D})$), subproblem selection ($\text{pop}(\mathcal{Q})$) and branching ($\text{split}(\mathcal{D}, \mathbf{y})$). The bounds computed by the bounding procedure are assumed to become tighter as the feasible set shrink. Especially when the size of a feasible set shrinks to a point, the difference between upper and lower bounds needs to converge to zero. It is shown in Proposition 1 that our SDP bounding approach satisfies these assumptions. A comprehensive study of branching and subproblem selection strategies can be found in [86, 87, 88, 89], in which a number of sophisticated rules are investigated. In this work, we adopt simple rules for branching and subproblem selection and rely on the tight bounding procedure to improve the performance.

3.2 Subproblem Selection: $\mathcal{D} = \text{pop}(\mathcal{Q})$

All unresolved subproblems (represented by their respective feasible sets) are stored in the priority queue \mathcal{Q} . At each iteration of B&B, the subproblem with the highest priority will be selected and split. In our implementation, subproblems are sorted according to the quality of their lower-bounds, and the subproblem with the worst lower-bound will be split firstly. In other words,

$$\mathcal{D} = \text{pop}(\mathcal{Q}) = \arg \min_{\mathcal{D}' \in \mathcal{Q}} \text{lb}(\mathcal{D}'), \quad (8)$$

Algorithm 1 Semidefinite Branch-and-Cut (SDBC)

Input: the original energy minimization problem $\min_{\mathbf{x} \in \mathcal{Z}^n} E(\mathbf{x})$.

Initialization: global upper-/lower-bounds $\text{gub} = +\infty$, $\text{glb} = -\infty$; priority queue $\mathcal{Q} = \{\mathcal{Z}^n\}$.

while $\mathcal{Q} \neq \emptyset$ **do**

Step 1 Subproblem Selection: $\mathcal{D} = \text{pop}(\mathcal{Q})$.

Step 2 Bounding:

 - $[\text{ub}, \text{lb}, \mathbf{x}, \mathbf{y}] = \text{bound}(\mathcal{D})$ (see Algorithm. 3);

Step 3 Update global bounds:

 - $\text{glb} = \min(\text{lb}, \min_{\mathcal{D}' \in \mathcal{Q}} \text{lb}(\mathcal{D}'))$;

 - **if** $\text{gub} > \text{ub}$ **then** $\text{gub} = \text{ub}$, $\mathbf{x}^* = \mathbf{x}$;

Step 4 Pruning: $\mathcal{Q} \leftarrow \mathcal{Q} \setminus \{\mathcal{D}' \mid \mathcal{D}' \in \mathcal{Q}, \text{lb}(\mathcal{D}') \geq \text{gub}\}$.

Step 5 Branching:

if $\text{lb} < \text{gub}$ **then**

 - $[\mathcal{D}_1, \mathcal{D}_2] = \text{split}(\mathcal{D}, \mathbf{y})$ (see Algorithm 2);

 - $\mathcal{Q} \leftarrow \{\mathcal{Q}, \mathcal{D}_1, \mathcal{D}_2\}$.

Output: \mathbf{x}^* .

Algorithm 2 The branching procedure, $[\mathcal{D}_1, \mathcal{D}_2] = \text{split}(\mathcal{D}, \mathbf{y})$.

Input: The search space $\mathcal{D} := \mathcal{Z}_1^{\mathcal{D}} \times \mathcal{Z}_2^{\mathcal{D}} \times \dots \times \mathcal{Z}_{|\mathcal{V}|}^{\mathcal{D}}$, and the singleton pseudomarginals $\{y_{p,i}\}_{i \in \mathcal{Z}_p^{\mathcal{D}}}, \forall p \in \mathcal{V}$.

Step 1 Node selection: $\varphi = \arg \min_{p \in \mathcal{V}, |\mathcal{Z}_p^{\mathcal{D}}| > 1} (\max_{i \in \mathcal{Z}_p^{\mathcal{D}}} y_{p,i})$.

Step 2 State ordering: obtain the state order $i_1, i_2, \dots, i_{|\mathcal{Z}_\varphi^{\mathcal{D}}|}$ such that $y_{\varphi, i_1} \geq y_{\varphi, i_2} \geq \dots \geq y_{\varphi, i_{|\mathcal{Z}_\varphi^{\mathcal{D}}|}}$.

Step 3 Branching: separate $\mathcal{Z}_\varphi^{\mathcal{D}}$ into two disjointed subsets $\mathcal{Z}_{\varphi,1}^{\mathcal{D}} = \{i_1, \dots, i_{\lfloor |\mathcal{Z}_\varphi^{\mathcal{D}}|/2 \rfloor}\}$ and $\mathcal{Z}_{\varphi,2}^{\mathcal{D}} = \{i_{\lfloor |\mathcal{Z}_\varphi^{\mathcal{D}}|/2 \rfloor + 1}, \dots, i_{|\mathcal{Z}_\varphi^{\mathcal{D}}|}\}$.

Output: $\mathcal{D}_1 = \mathcal{Z}_1^{\mathcal{D}} \times \dots \times \mathcal{Z}_{\varphi,1}^{\mathcal{D}} \times \dots \times \mathcal{Z}_{|\mathcal{V}|}^{\mathcal{D}}$ and $\mathcal{D}_2 = \mathcal{Z}_1^{\mathcal{D}} \times \dots \times \mathcal{Z}_{\varphi,2}^{\mathcal{D}} \times \dots \times \mathcal{Z}_{|\mathcal{V}|}^{\mathcal{D}}$.

where $\mathcal{D} := \mathcal{Z}_1^{\mathcal{D}} \times \mathcal{Z}_2^{\mathcal{D}} \times \dots \times \mathcal{Z}_{|\mathcal{V}|}^{\mathcal{D}}$ and $\mathcal{Z}_p^{\mathcal{D}}$ denotes the set of possible states for each node $p \in \mathcal{V}$ within the feasible set \mathcal{D} . Note that a subset of elements of \mathbf{y} are enforced to be 0s or 1s in the subproblem with respect to \mathcal{D} :

$$y_{p,i} = 0, \quad \text{if } i \in \{\mathcal{Z} \setminus \mathcal{Z}_p^{\mathcal{D}}\}, \forall p \in \mathcal{V}, \quad (9a)$$

$$y_{p,i} = 1, \quad \text{if } i \in \mathcal{Z}_p^{\mathcal{D}} \text{ and } |\mathcal{Z}_p^{\mathcal{D}}| = 1, \forall p \in \mathcal{V}. \quad (9b)$$

3.3 Branching: $[\mathcal{D}_1, \mathcal{D}_2] = \text{split}(\mathcal{D}, \mathbf{y})$

At the branching step, the input feasible set \mathcal{D} is further split into two disjointed sets \mathcal{D}_1 and \mathcal{D}_2 , by selecting a node and split its state space in halves.

Node Selection We employ the “difficult first” strategy, which selects the node that is the most difficult to discretize. Remind that $\{y_{p,i}\}_{i \in \mathcal{Z}_p^{\mathcal{D}}}$ can be considered as a singleton pseudomarginal for each node $p \in \mathcal{V}$, if it resides in the local marginal polytope (4). If this distribution does not concentrate on a single point (in other words, spreads evenly over all states), then the corresponding node is considered to be difficult to discretize. Accordingly, a simple strategy is adopted for node selection in this work:

$$\varphi = \arg \min_{p \in \mathcal{V}, |\mathcal{Z}_p^{\mathcal{D}}| > 1} (\max_{i \in \mathcal{Z}_p^{\mathcal{D}}} y_{p,i}). \quad (10)$$

States Ordering As in [71], the possible states of the selected node φ , i.e. $i_1, i_2, \dots, i_{|\mathcal{Z}_\varphi^\mathcal{D}|}$, are ordered dynamically in each iteration of branching, such that $y_{\varphi, i_1} \geq y_{\varphi, i_2} \geq \dots \geq y_{\varphi, i_{|\mathcal{Z}_\varphi^\mathcal{D}|}}$. Then the most and the least probable states are separated into two sets: $\mathcal{Z}_{\varphi,1}^\mathcal{D} = \{i_1, \dots, i_{\lfloor |\mathcal{Z}_\varphi^\mathcal{D}|/2 \rfloor}\}$ and $\mathcal{Z}_{\varphi,2}^\mathcal{D} = \{i_{\lfloor |\mathcal{Z}_\varphi^\mathcal{D}|/2 \rfloor + 1}, \dots, i_{|\mathcal{Z}_\varphi^\mathcal{D}|}\}$.

Now the last component in the B&C approach (Algorithm 1) to be studied is the bounding procedure, $\text{bound}(\mathcal{D})$, which we discuss in the following section.

4 An Efficient SDP Bounding Procedure with Cutting-plane

In this section, we consider the following SDP relaxation to the original energy minimization problem (1):

$$\min_{\mathbf{y}, \mathbf{Y}} E(\mathbf{y}, \mathbf{Y}) \quad (11a)$$

$$\text{s.t. (12), (13), (14), (15), (16), (17), (18), (19),} \quad (11b)$$

$$\Omega(\mathbf{y}, \mathbf{Y}) \succeq \mathbf{0}, \quad (11c)$$

where (11b) corresponds to the additional linear constraints to be discussed in Section 4.1. The above formulation can be easily extended to any subproblem \mathcal{D} in Algorithm 1, by reparameterization or enforcing additional constraints (9). A main observation is that the linear constraints arising from standard/high-order LP relaxation can be used to significantly tighten the following naive SDP relaxation: However, the resulted large-scale SDP optimization problem, which in particular has a large number of linear constraints, render standard interior-point SDP solvers inapplicable. A scalable SDP solver combined with cutting-plane is proposed to solve such SDP optimization in Section 4.2 and 4.3. Then the complexity of our method is discussed in Section 4.4 and several speeding-up strategies are proposed in Section 4.5.

4.1 Linear Constraints that Tighten the Semidefinite Outer-Bound

Besides the p.s.d. constraint, several classes of linear constraints arising from LP relaxation can be added to SDP relaxation to tighten the bound.

0/1-Integer Constraints These constraints arise from the integer constraints $y_{p,i} \in \{0, 1\}$, which is equivalent to $y_{p,i} = y_{p,i}^2$, for each node $p \in \mathcal{V}$ and state $i \in \mathcal{Z}$:

$$Y_{pi,pi} = y_{p,i}, \quad \forall p \in \mathcal{V}, i \in \mathcal{Z}. \quad (12)$$

There are nh such constraints for a graph with n nodes and h states per node.

Local Normalization Constraints Because only one state $i \in \mathcal{Z}$ can be assigned to each node $p \in \mathcal{V}$, we have:

$$\sum_{i \in \mathcal{Z}} y_{p,i} = 1, \quad \forall p \in \mathcal{V}, \quad (13)$$

There are n local normalization constraints in total. Note that for any $\{\mathbf{y}, \mathbf{Y}\} \in \mathcal{P}(\mathcal{G}) \cap (12) \cap (13)$, if $\text{rank}(\Omega(\mathbf{y}, \mathbf{Y})) = 1$, then $\{\mathbf{y}, \mathbf{Y}\}$ is a vertex of $\mathcal{M}(\mathcal{G}, \mathcal{Z})$.

Non-Negativity Constraints Deriving from $Y_{pi,qi} = y_{p,i} \cdot y_{q,i}$ and $\mathbf{y} \in \{0, 1\}^{nh}$, the following constraints holds:

$$Y_{pi,qj} \geq 0, \quad \forall (p, q) \in \mathcal{E}, i, j \in \mathcal{Z}. \quad (14)$$

There are $|\mathcal{E}|h^2$ such constraints for a graph with $|\mathcal{E}|$ edges. The number becomes $n(n-1)h^2/2$ for a fully-connected graph. Huang *et al.* [50] have shown that without the non-negativity constraints, SDP relaxation is loose for submodular functions.

Edge Marginalization Constraints For each edge $(p, q) \in \mathcal{E}$ and $i \in \mathcal{Z}$, we have $\sum_{j \in \mathcal{Z}} Y_{pi,qj} = y_{pi} \cdot \sum_{j \in \mathcal{Z}} y_{qj}$. Therefore the following constraints can be derived:

$$\sum_{j \in \mathcal{Z}} Y_{pi,qj} = y_{pi}, \quad \forall (p, q) \in \mathcal{E}, i \in \mathcal{Z}. \quad (15)$$

There are in total $2|\mathcal{E}|h$ such constraints ($n(n-1)h$ for fully-connected graphs). Note that local normalization (13), non-negativity (14) and marginalization (15) constraints force $\{\mathbf{y}, \mathbf{Y}\}$ to lie in local marginal polytope $\mathcal{M}_L(\mathcal{G}, \mathcal{Z})$.

Gangster Operators Because $Y_{pi,pj} = y_{p,i} \cdot y_{p,j}$ and each node $p \in \mathcal{V}$ can only be assigned with one state $i \in \mathcal{Z}$, at least one of $y_{p,i}$ and $y_{p,j}$ should be zero. Then we have the following constraints:

$$Y_{pi,pj} = 0, \quad \forall p \in \mathcal{V}, \forall i \neq j, i, j \in \mathcal{Z}. \quad (16)$$

These constraints are referred to as *gangster operators*, as they shoot holes (zeros) in \mathbf{Y} . Considering matrix symmetry, there are $nh(h-1)/2$ such constraints in total. Schellewald and Schnorr [46] have used these constraints for SDP approaches to subgraph matching.

Note that the above linear constraints are directly defined on the marginal polytope. In this work, we also consider the following a few classes of linear constraints defined on the cut polytope or equivalently binary marginal polytope, which can be extended to $\mathcal{M}(\mathcal{G}, \mathcal{Z})$ by projecting a non-binary graph to a binary graph. Comprehensive studies of linear constraints for cut polytope or binary marginal polytope can be found in [90, 91, 92] and the references therein.

Triangular Inequalities Considering three binary variables π_p, π_q, π_r and define $\delta_{p,q} := \mathbf{1}[\pi_p \neq \pi_q]$, then we have:

$$\delta_{p,q} + \delta_{q,r} + \delta_{p,r} \leq 2, \quad (17a)$$

$$\delta_{p,q} - \delta_{q,r} - \delta_{p,r} \leq 0, \quad (17b)$$

$$-\delta_{p,q} + \delta_{q,r} - \delta_{p,r} \leq 0, \quad (17c)$$

$$-\delta_{p,q} - \delta_{q,r} + \delta_{p,r} \leq 0, \quad (17d)$$

which are facet defining for a cut polytope [90]. The number of triangular inequalities is cubic in the number of binary variables.

Cycle Inequalities They arise from the fact that there must be an even number of edge-cuts for any cycle of variables. For a cycle C and any $F \subseteq C$ such that $|F|$ is odd, the

inequality can be expressed as:

$$\sum_{(p,q) \in C \setminus F} \delta_{p,q} + \sum_{(p,q) \in F} (1 - \delta_{p,q}) \geq 1. \quad (18)$$

The triangular inequalities (17b) (17c), (17d) can be considered as special cases of cycle inequalities. A cycle inequality is facet defining for a cut polytope if and only if the cycle is chordless [90]. There are an exponential number of cycle inequalities in a graph. However, the most violated one can be found in $\mathcal{O}(n^3 h^3)$ time for a dense graph using the separation procedure shown in [90, 93].

Odd-wheel Inequalities Consider a cycle C of length $s = 2t + 3$, where the integer $t \geq 0$. $V(C)$ denotes the node set of C and π_r refers to a center node not in $V(C)$. Then the following inequality

$$\sum_{(p,q) \in C} \delta_{p,q} - \sum_{p \in V(C)} \delta_{p,r} \leq \frac{|C| - 1}{2}, \quad (19)$$

is called ‘ s -wheel inequality’ or simply ‘odd-wheel inequality’. This class of constraints is not facet-defining for a cut polytope, but it is facet defining for a multicut polytope [91, 92]. There are also an exponential number of the above constraints with respect to the number of binary variables. For a solution satisfying all cycle inequalities, the separation procedure shown in [93] can find the most violated odd-wheel inequality using $\mathcal{O}(n^4 h^4)$ time for a dense graph.

We adopt the method proposed in [42, 52] for projecting non-binary graphs to binary graphs. Firstly, π_p^s is defined as a binary variable separating the state space of node p into two disjoint non-empty sets \mathcal{Z}_s and $\mathcal{Z} \setminus \mathcal{Z}_s$, that is $\pi_p^s = 1$, if $x_p \in \mathcal{Z}_s$, and 0 otherwise. $\pi_p = \{\pi_p^1, \pi_p^2, \dots\}$ represents a collection of partitions with respect to node p . Then the non-binary graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ can be projected to a binary graph $\mathcal{G}_\pi = (\mathcal{V}_\pi, \mathcal{E}_\pi)$, where

$$\mathcal{V}_\pi = \{\pi_p^s \mid p \in \mathcal{V}, s \leq |\pi_p|\}, \quad (20a)$$

$$\mathcal{E}_\pi = \{(\pi_q^s, \pi_p^t) \mid (p, q) \in \mathcal{E}, s \leq |\pi_p|, t \leq |\pi_q|\}. \quad (20b)$$

Note that $\delta_{\pi_q^s, \pi_p^t} := \mathbf{1}[\pi_p^s \neq \pi_q^t]$ can be expressed in terms of \mathbf{Y} in (5), that is $\delta_{\pi_q^s, \pi_p^t} = \sum_{\{i \in \mathcal{Z}_s, j \notin \mathcal{Z}_t\} \cap \{i \notin \mathcal{Z}_s, j \in \mathcal{Z}_t\}} Y_{pi, qj}$. In this paper, we define $\mathcal{Z}_s = \{s\}, \forall s = 1, 2, \dots, h$, such that each h -state variable x_p is projected to h binary variables.

Note that the projected inequalities are not necessarily facet-defining for the non-binary graph, although some of them are facet-defining on the projected binary graph.

4.2 Initial Bounding Using SDCut

The feasible set of the SDP relaxation (11) is the intersection of semidefinite outer-bound $\mathcal{P}(\mathcal{G})$ and the polyhedral outer-bound produced by standard/high-order LP relaxation. There are several difficulties in optimizing the SDP problem (11). First of all, it is computationally inefficient to solve SDP problems using classic interior-point methods, which

Algorithm 3 SDP Bounding Procedure ($[\text{ub}, \text{lb}, \mathbf{x}, \mathbf{y}] = \text{bound}(\mathcal{D})$)

Input: The energy minimization problem $\min_{\mathbf{x} \in \mathcal{D}} E(\mathbf{x})$, where \mathcal{D} is a subset of the space of potential label assignment \mathcal{Z}^n ; Ψ is the working set of constraints. Ψ_{addi} is the pool of additional constraints (14), (15), (17), (18) and (19) to be added in cutting-plane. Parameters $\gamma > 0, \delta > 1$; Maximum iterations $K_{\text{init}}, K_{\text{inner}}$ and K_{outer} . The initial set of constraints ψ_{init} and the corresponding dual variable \mathbf{u}_{init} .

Step 1 Initial Bounding:

- Model reduction (Optional): $\mathcal{D} \leftarrow \text{QPBO}(\mathcal{D})$;
- Initialize the working set and dual variable: $\Psi \leftarrow \psi_{\text{init}}, \mathbf{u} \leftarrow \mathbf{u}_{\text{init}}$.
- Construct the dual problem (21) with respect to \mathcal{D}, Ψ and γ .

for $k = 1, 2, \dots, K_{\text{init}}$ **do**

- Dual variable update: $\mathbf{u} \leftarrow \mathbf{u} + \rho \mathbf{H} \nabla d_\gamma(\mathbf{u})$.
- Lower-bound update: $\text{lb} \leftarrow d_\gamma(\mathbf{u})$.
- Goto **Step 3** if any stop condition is met (see Section 4.5).

- Compute the primal variable: $\Omega(\mathbf{y}, \mathbf{Y}) = \gamma \Pi_{\mathcal{S}_+^{nh+1}}(\mathbf{C}(\mathbf{u}))$.

Step 2 Cutting-Plane:

for $k_1 = 1, 2, \dots, K_{\text{outer}}$ **do**

- $\psi_{\text{in}} \leftarrow$ inactive constraints selected from Ψ .
- $\psi_{\text{vi}} \leftarrow$ constraints violated by $\Omega(\mathbf{y}, \mathbf{Y})$ selected from Ψ_{addi} .
- $\Psi \leftarrow \{\Psi \setminus \psi_{\text{in}}\} \cup \psi_{\text{vi}}, \gamma \leftarrow \delta \cdot \gamma$.
- Construct the dual problem (21) with respect to \mathcal{D}, Ψ and γ .
- Update the dual variable: $\mathbf{u} \leftarrow \{\mathbf{u}, \mathbf{0}\}$, where $\mathbf{0}$ corresponds to the dual variables with respect to ψ .

for $k_2 = 1, 2, \dots, K_{\text{inner}}$ **do**

- Dual variable update: $\mathbf{u} \leftarrow \mathbf{u} + \rho \mathbf{H} \nabla d_\gamma(\mathbf{u})$.
- Lower-bound update: $\text{lb} \leftarrow d_\gamma(\mathbf{u})$.

- Goto **Step 3** if any stop condition is met.

- Compute the primal variable: $\Omega(\mathbf{y}, \mathbf{Y}) = \gamma \Pi_{\mathcal{S}_+^{nh+1}}(\mathbf{C}(\mathbf{u}))$.

Step 3 Rounding:

- Generate a discrete solution \mathbf{x} by rounding \mathbf{y} (see Equation (7)).
- Improve \mathbf{x} using RICM (see Section 4.5).
- $\text{ub} \leftarrow E(\mathbf{x})$.

Output: $\text{ub}, \text{lb}, \mathbf{x}$ and \mathbf{y} .

has the computational complexity of $\mathcal{O}(m^3 + mn^3 h^3 + m^2 n^2 h^2)$ and memory requirement of $\mathcal{O}(m^2 + n^2 h^2)$, where m is the number of linear constraints. Secondly, the linear constraints involved in (11) are very complex. It requires $\mathcal{O}(n^2 h^2)$ constraints to describe the local marginal polytope for a fully-connected graph, which makes the computational complexity of interior-point methods increase to $\mathcal{O}(n^6 h^6)$. The total number of cycle inequalities grows even exponentially with the graph size. To this end, we propose to adapt the SDCut method proposed in our prior work [49] to optimize (11) and use cutting-plane to cope with the large number of linear constraints.

In the following, $\mathbf{A}, \mathbf{B}_i \in \mathcal{S}^{nh+1}, i = 1, 2, \dots, m, \mathbf{b} \in \mathbb{R}^m$ are defined such that $\langle \Omega(\mathbf{y}, \mathbf{Y}), \mathbf{A} \rangle = E(\mathbf{y}, \mathbf{Y})$ and linear constraints in (11b) are also expressed in the form of matrix inner-products: $\langle \Omega(\mathbf{y}, \mathbf{Y}), \mathbf{B}_i \rangle = b_i, \forall i \in \mathcal{J}_{\text{eq}}, \langle \Omega(\mathbf{y}, \mathbf{Y}), \mathbf{B}_i \rangle \leq b_i, \forall i \in \mathcal{J}_{\text{in}}$, where $\mathcal{J}_{\text{eq}}, \mathcal{J}_{\text{in}}$ refer to non-overlapping indexes for linear equality and inequality constraints respectively. $m = |\mathcal{J}_{\text{eq}}| + |\mathcal{J}_{\text{in}}|$.

SDCut [49] is proposed to approximately solve a subclass of SDP problems in which the trace of p.s.d. matrix variables is fixed. Instead of minimizing the linear objec-

tive function $\langle \Omega(\mathbf{y}, \mathbf{Y}), \mathbf{A} \rangle$, SDCut minimizes the perturbed function $\langle \Omega(\mathbf{y}, \mathbf{Y}), \mathbf{A} \rangle + \frac{1}{2\gamma} (\|\Omega(\mathbf{y}, \mathbf{Y})\|_F^2 - \eta^2)$ over the constraints shown in (11b), where $\eta := \text{trace}(\Omega(\mathbf{y}, \mathbf{Y})) = n+1$ (η can be discarded since it is a constant) in our case and $\gamma > 0$. It is shown in [49] that with an appropriate γ , the solution of the perturbed problem is “close” to that of the original SDP problem.

Furthermore, by introducing the perturbed objective, the Lagrangian dual of SDCut can be simplified to:

$$\max_{\mathbf{u} \in \mathbb{R}^m} d_\gamma(\mathbf{u}) := -\frac{\gamma}{2} \|\Pi_{\mathcal{S}_+^{nh+1}}(\mathbf{C}(\mathbf{u}))\|_F^2 - \mathbf{u}^\top \mathbf{b} - \frac{\eta^2}{2\gamma}, \quad (21a)$$

$$\text{s.t. } u_i \geq 0, i \in \mathcal{I}_{in}, \quad (21b)$$

where $\mathbf{C}(\mathbf{u}) = -\mathbf{A} - \sum_{i=1}^m u_i \mathbf{B}_i$. The matrix-valued function $\Pi_{\mathcal{S}_+^{nh+1}}$ is defined as

$$\Pi_{\mathcal{S}_+^{nh+1}}(\mathbf{X}) := \mathbf{P} \text{diag}(\max(0, \boldsymbol{\lambda})) \mathbf{P}^\top,$$

and $\mathbf{X} = \mathbf{P} \text{diag}(\boldsymbol{\lambda}) \mathbf{P}^\top$ is the eigen-decomposition of the matrix $\mathbf{X} \in \mathcal{S}^n$, where $\boldsymbol{\lambda}$ denotes the vector of eigenvalues and \mathbf{P} contains the corresponding (column) eigenvectors. Note that there is no p.s.d. constraint in (21) any more. Based on K.K.T. conditions, we have

$$\Omega(\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*) = \gamma \Pi_{\mathcal{S}_+^{nh+1}}(\mathbf{C}(\mathbf{u}_\gamma^*)), \quad (22)$$

where $\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*$ and \mathbf{u}_γ^* are the optimal solutions to the corresponding primal and dual problems with respect to γ respectively.

We demonstrate in Appendix 7.1 that the dual formulation of SDCut (21) is actually equivalent to a penalty approach to the Lagrangian dual of the original SDP relaxation (11). For an sufficiently large parameter γ , the simplified dual (21) can be arbitrarily close to the original SDP relaxation (11), which however results in numerical problems. In many cases, it is sufficient to restrict γ to a large value and obtain an approximate solution. Note that how to choose the value of γ in practice to make a good trade-off between bound quality and convergence speed is still an open question. A potential strategy is to adjust γ adaptively at each iteration of descent step.

The dual objective function $d_\gamma(\cdot)$ is continuously differentiable but not necessarily twice differentiable, and its gradient is given by

$$\nabla d_\gamma(\mathbf{u}) = -\gamma \Phi \left[\Pi_{\mathcal{S}_+^{nh+1}}(\mathbf{C}(\mathbf{u})) \right] - \mathbf{b}, \quad (23)$$

where $\Phi : \mathcal{S}^{nh+1} \rightarrow \mathbb{R}^m$ refers to the linear transformation $\Phi[\Omega] := [\langle \mathbf{B}_1, \Omega \rangle, \dots, \langle \mathbf{B}_m, \Omega \rangle]^\top$. Consequently, we can solve the dual (21) using quasi-Newton algorithms as they only require first-order derivatives during optimization. At each descent step, quasi-Newton methods update the dual variable \mathbf{u} as follows:

$$\mathbf{u} \leftarrow \mathbf{u} + \rho \mathbf{H} \nabla d_\gamma(\mathbf{u}), \quad (24)$$

where \mathbf{H} is the approximated inverse of Hessian matrix updated by successive gradient vectors and $1 \geq \rho \geq 0$ is the step size decided by line-search.

Another important property of the simplified dual problem (21) is that for any feasible dual variable \mathbf{u} , the dual objective function value $d_\gamma(\mathbf{u})$ is a lower-bound to the minimum energy $\min_{\mathbf{x} \in \mathcal{Z}^n} E(\mathbf{x})$ [49]. This is beneficial for accelerating the B&B methods: the bounding procedure can be stopped once the computed lower-bound is not lower than the current global upper-bound, and the corresponding subproblem can be pruned.

In the sequel, $d_\gamma^*(\mathcal{D})$ is defined as the optimal value of the dual problem (21) constructed with respect to γ and $\mathcal{D} \subseteq \mathcal{Z}^n$. In this paper, we further prove the following results for the tightness of the lower-bound yielded by $d_\gamma^*(\mathcal{D})$ (see Appendix 7.2):

Proposition 1 *The following results hold: (i) $\forall \gamma > 0, \forall \mathcal{D}_1 \subseteq \mathcal{D}_2 \subseteq \mathcal{Z}^n$, we have $d_\gamma^*(\mathcal{D}_1) \geq d_\gamma^*(\mathcal{D}_2)$; (ii) $\forall \gamma > 0, \forall \mathcal{D} \subseteq \mathcal{Z}^n$ with $|\mathcal{D}| = 1$, we have $d_\gamma^*(\mathcal{D}) = \min_{\mathbf{x} \in \mathcal{D}} E(\mathbf{x})$.*

The first result shows that the lower-bound is improved as the search space \mathcal{D} shrinks, and the second means that the gap between the lower- and upper-bound $E(\mathbf{x})_{\mathbf{x} \in \mathcal{D}} - d_\gamma^*(\mathcal{D})$ converges to zero when the search space \mathcal{D} shrinks to a point (namely $|\mathcal{D}| = 1$). In particular, if $|\mathcal{D}| = 1$ and only contains the MAP solution, $\arg \min_{\mathbf{x} \in \mathcal{Z}^n} E(\mathbf{x})$, $d_\gamma^*(\mathcal{D})$ equals to the global minimum energy. These two properties are crucial when embedding the proposed bounding method into branch-and-bound.

Proposition 2 *(Optimality Guarantee of SDP Relaxation) For any $\gamma > 0$, suppose $\{\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*\}$ and \mathbf{u}_γ^* are the optimal/dual primal variable obtained by (22). If the rank of $(\Omega(\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*))$ is 1, then $\{\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*\}$ yields the exact MAP solution and $d_\gamma(\mathbf{u}_\gamma^*)$ is the minimum energy value.*

The above optimality guarantee of SDP relaxation can be considered as a counterpart to Theorem 1 for the standard LP relaxation. Both LP and SDP relaxations drop non-convex constraints and result in a simple convex relaxation: the integer constraint is relaxed by LP relaxation and SDP relaxation drops the rank-1 constraint. Note that Proposition 2 can be generalized to SDP relaxation with linear constraints (12), (13).

Computationally SDCut is significantly more scalable and efficient than general interior-point methods. It is employed here to solve the SDP relaxation with linear constraints in (29c). However, there are $\mathcal{O}(n^2 h^2)$ number of constraints (14), (15) and the exponential number of constraints (17), (18), (19). Adding these constraints directly to SDCut is still impractical. To this end, we proposed to find and add most violated constraints via cutting-plane, which is discussed in the next section.

4.3 Cutting-Plane: Searching and Adding Violated Constraints

In practice, we find that most of the constraints (14), (15), (17), (18) and (19) are redundant after the initial bounding procedure shown in Section 4.2, which means that most of them are already satisfied after imposing the p.s.d. constraints and the linear constraints (12), (13) and (16). Therefore, cutting-plane methods can be adopted to search violated constraints in (14), (15), (17), (18) and (19), and add them to SDP relaxation at each iteration. Adding these primal constraints is equivalent to adding new variables to the dual problem (21).

Violated non-negativity (14), edge marginalization (15) constraints and triangular inequalities (17) are found by enumeration. Most violated cycle inequalities (18) are found using the separation procedure [93]. When there is no violated cycle inequality (up to certain precision), violated odd-wheel inequalities (19) are searched and added using the separation procedure shown in [93].

The procedure of the proposed SDP bounding procedure, including initial bounding (Section 4.2) and cutting-plane (Section 4.3) are summarized in Algorithm 3. At each iteration of cutting-plane, the dual problem (21) is constructed based on three inputs: the search space \mathcal{D} (for reparameterization), linear constraints set Ψ and the penalty parameter γ .

4.4 Computational Complexity

L-BFGS-B [94] is used as the implementation of quasi-Newton algorithms for SDCut. At each iteration, L-BFGS-B itself requires only $\mathcal{O}(m)$ arithmetic operations and memory requirement, so SDCut scales linearly in the number of constraints m . The main computational burden of SDCut is the eigen-decomposition at each gradient-descent step to compute the dual objective function (21a) and its gradient (23). The eigen-solver used in this paper is the DSEVR routine in LAPACK (an implementation of Relatively Robust Representations), which has the computational complexity of $\mathcal{O}(n^3h^3)$ and the memory requirement of $\mathcal{O}(n^2h^2)$.

Finding the most violated non-negativity constraint (14) and edge marginalization constraint (15) requires $\mathcal{O}(n^2h^2)$ arithmetic operations, while finding violated cycle inequality needs $\mathcal{O}(n^3h^3)$ arithmetic operations. To save time, only a few center node π_r (see Equation (19)) is considered for odd-wheel inequalities in each cutting-plane iteration.

Therefore the overall computational complexity of the proposed SDP method is $\mathcal{O}(n^3h^3)$ at each descent iteration (m is less than $\mathcal{O}(n^3h^3)$, so it is omitted), and the corresponding memory requirement is $\mathcal{O}(n^2h^2)$. This result demonstrates that our approach is much more efficient than the

classic interior-point methods, which needs $\mathcal{O}(n^6h^6)$ arithmetic operations.

4.5 Strategies to Speed Up

Although the proposed SDP approximation method already provides a tight lower-bound efficiently, we proposed a few techniques to further speed up the bounding procedure of B&B.

Pre-Processing (Model Reduction) Several methods have been proposed to identify persistency for binary models [9], Potts models [95, 96] and general multi-label models [95, 97, 98, 99, 100, 101]. Persistencies refer to states of variables which are proved to be non-optimal or belong to at least one optimal solution. Previous work has employed persistency to reduce the model size before applying approximate [102] and exact [103] solvers. For binary models, all variables assigned as integers by QPBO [9] are persistent.

Without affecting the property of global convergence of B&B, QPBO is performed in each bounding procedure, and then the SDP bounding procedure is applied on the reduced model. In general, the computation time of QPBO is negligible compared to that of the SDP procedure. Because the size of the resulted SDP is smaller, the total computational time is reduced as well. Note that this partial optimality also holds for TRW message passing and other methods equivalent to the standard LP relaxation for binary graphs. We choose QPBO here for its fast computation.

Stop Conditions As discussed in Section 4.2, the objective value of (21), calculated at each descent step, can be cast as a lower-bound to the minimum energy. For faster computation, the bounding process is stopped before convergence if any of the following conditions is met (see Algorithm 3):

1. The lower-bound lb is already greater than the global upper-bound gub, and thus this subproblem is pruned.
2. The improvement of the lower-bound lb between consecutive steps is smaller than the specified tolerance.

Warm Start Another strategy used in the bounding procedure to speed up the computation is “warm start”, which is performed in two levels: between different iterations of cutting-planes and bounding procedures.

1. In an iteration of cutting-plane, the dual variables are initialized by the results of the previous iteration. For the new constraints, the corresponding dual variables are filled with zeros.
2. The final working set of active constraints and the corresponding dual variables of the first bounding iteration are used to initialize those of subsequent bounding iterations (ψ_{init} and \mathbf{u}_{init} in Algorithm 3), which speed up the convergence speed of L-BFGS-B. The initial working set ψ_{init} of the first bounding iteration is set to con-

tain constraints (12), (13) and (16), and \mathbf{u}_{init} is set to zeros.

Removing Inactive Constraints One disadvantage of cutting plane is that the size of the subproblems (the number of primal linear constraints, and correspondingly the number of dual variables) keeps increasing from iteration to iteration,

Since traditionally new constraints are added in each round but are never removed. The strategy of pruning inactive constraints has been employed to address this issue [104, 105, 106] in cutting-plane methods. In this work, we use a simple rule to decide which constraint should be dropped from the working set: those corresponding to zero dual variables and not violated by the current primal solution. Note that those constraints discarded by the current iteration may be re-activated in subsequent cutting-plane iterations.

Post-Processing (Stochastic Local Search)

When the dual optimization is terminated, an approximate solution $\mathbf{x}_{\text{appx}} \in \mathcal{Z}^n$ is calculated by rounding and an upper-bound $E(\mathbf{x}_{\text{appx}})$ to the minimum energy $\min_{\mathbf{x} \in \mathcal{Z}^n} E(\mathbf{x})$ is obtained. Local (greedy) search methods can be applied to further tighten the upper-bound. What we use here is ‘repeated iterated conditional modes’ (refer to as RICM):

1. Run Iterated Conditional Modes (ICM) [83] with \mathbf{x}_{appx} as the initial point to find a local optimal solution $\mathbf{x}_{\text{local}}$. Compute the upper-bound $E(\mathbf{x}_{\text{local}})$.
2. Randomly flip a portion (around 5%) of assignments in the initial solution \mathbf{x}_{appx} .
3. Repeat the steps 1 and 2 several times and select the local optimal solution $\mathbf{x}_{\text{local}}$ with the lowest upper-bound as the new approximate solution.

Note that ICM can be considered as a greedy coordinate descent method. The solution quality is guaranteed to be improved or at least not worse than the initial point. Furthermore, as ICM depends highly on the initialization, we repeat ICM several times to overcome this problem.

5 Experiments

The proposed method is evaluated in two forms: one with cutting-planes (refer to as SDBC) and one without (refer to as SDBB). Both of the SDP bounding methods in SDBB and SDBC solve the SDP relaxation (11) approximately based on the simplified dual formulation (21). The only difference is that SDBC imposes all linear constraints listed in (11b)¹, while SDBB only considers constraints (12), (13) and (16). The following algorithms for MAP-MRF inference are compared against ours.

¹ In the following experiments, we find that odd-wheel inequalities are only effective on the modularity clustering models, therefore this class of constraints is not considered for other models

1) As a baseline, the lower-bound produced by TRWS [18] is compared with SDBC.

2) Sontag *et al.* [39, 42] proposed to tighten the LP relaxation to MAP problems by incrementally introducing cluster consistency constraints, for example, triplet clusters [39] (refer to as MPLP-CP-v1) or long cycles [42] (refer to as MPLP-CP-v2). The code is obtained from the author’s website. At each round, 50 additional triplets are added to MPLP-CP-v1 and 50 ~ 200 triplets plus 50 ~ 200 cycles are added to MPLP-CP-v2. MPLP-CP-v1 and MPLP-CP-v2 provide lower-bounds to the global optimal energy value, which can be used to validate the exactness of a solution. For SDP relaxation methods, we compared our method with SDPT3, which represents one of the state-of-the-art implementations of interior-point methods.

3) Our method are also compared with methods based on graph cuts [8, 107] for binary non-submodular pairwise energies. QPBO [8] can be used to exactly solve problems with a small fraction of non-submodular terms. We also use QPBO to reduce the model before applying other methods. Local submodular approximation (LSA) [107] iteratively approximates non-submodular energies non-linearly based on trust region (refer to as LSA-TR) or auxiliary function principles.

4) Several B&B methods for energy minimization are also evaluated. Sun *et al.* [71] proposed a B&B method (refer to as MPLP-BB) using MPLP to upper-bound the MAP value at each branch. The dual variables of MPLP at each branch are fixed, and a data structure, called range maximum query, is pre-constructed to speed up the bound calculation significantly. The parameters of MPLP-BB are set to the default values, except that the number of initial MPLP iterations is set to 1000.

DAOOPT [72] implements the And/Or Branch and Bound paradigm, which also uses mini-bucket heuristic for pruning search space and limited discrepancy search (LDS) to find an initial solution quickly. It achieves the best performance in the PIC2011 Challenge [108]. In this paper, the memory limit of DAOOPT is set to 8GB and the discrepancy limit of LDS is set to 2. The commercial integer linear program solver, IBM CPLEX [109] (denoted by ILP), is also evaluated in the experiments. Some specialized algorithms [110, 111, 112, 113] are also investigated for some particular data sets.

To perform a fair comparison, all the implementations of compared methods are obtained from the author’s website or from the OpenGM2 toolbox [114] and tested on one core of a 2.7GHz Intel CPU with 4GB free memory, except for DAOOPT (10GB) and ILP (40GB). Note that some experimental results of competitive methods are obtained from [114], which is evaluated on a faster Intel CPU (3.2GHz). All the methods are performed with the same procedures of pre-processing (model reduction by QPBO) and post-processing




Fig. 1: Lower-bounds given by different algorithms on the model “deer_rescaled_0034.K15.F100” (60 variables, 15 states per variable) in PIC2011. SDBC is performed without branching (SDBC-NoBranch) or without both branching and cutting-plane (SDBB-NoBranch). The lower-bounds provided by LP-based methods (TRWS and MPLP-CP) are worse than SDP-based methods (SDPT3 and SDBC). SDBB-NoBranch produces a lower-bound very close to the interior-point method SDPT3. For SDBC itself, the additional constraints added through cutting-plane gives a 5% increase in lower-bound.

(stochastic local research by ICM). The primary results (Table 2, 4, 5 and 6) are shown in a format similar to [1], which demonstrates the upper-bound (energy value of the final integer solution) and lower-bound (to the optimal energy value), averaged over all instances of each problem. For an algorithm, the number of instances for which it achieves the best upper-/lower-bound among all algorithms (denoted as #Best-ub and #Best-lb respectively) and the number of instances to which it gives the exact solution (denoted as #Exact) are also demonstrated for better comparison. As in [1], a solution is considered to be exact if the difference between the lower-bound and upper-bound is less than 10^{-5} in terms of absolute value or less than 10^{-8} in terms of relative value.

The implementation of our algorithm is written in matlab and the implementation of other methods are mainly implemented in C/C++. We would expect a further improvement in the speed of our method if it were implemented in optimized C/C++.

5.1 Performance Analysis

Relaxation Tightness The SDP bounding procedure embedded in SDBC is compared with state-of-the-art SDP/LP relaxation methods. SDPT3 [115], an implementation of interior point methods, is used to accurately solve the SDP relaxation with linear constraints (12), (13) and (16). Note that the other classes of tightening constraints are difficult to be imposed with SDPT3, due to the poor scalability of interior-point methods in the number of linear constraints. MPLP-CP-v1 and MPLP-CP-v2 tighten the standard LP relaxation

(over local marginal polytope $\mathcal{M}_L(\mathcal{G}, \mathcal{Z})$) by incrementally adding high-order consistency constraints.

In Fig. 1, the performance of these relaxation methods are compared on the lower-bounds achieved on a model in the PIC2011, “deer_rescaled_0034.K15.F100”, which is a fully-connected graph with 60 nodes and 15 states per node. In this experiment, SDBC is performed without branching, which means there is only one bounding procedure. Furthermore, to show the effectiveness of cutting-plane, SDBC is performed in two forms: with (SDBC-NoBranch) or without cutting-plane (SDBB-NoBranch).

From the results in Fig. 1, we can see that:

1. SDBB-NoBranch achieves a final lower-bound very similar to that produced by SDPT3 (the difference is only 0.1%), but is significantly faster than SDPT3. Note that SDBB-NoBranch and SDPT3 solve the same SDP problem, but approximately and exactly respectively. As more linear constraints are imposed, SDBC-NoBranch achieves better lower-bound than SDBB-NoBranch and SDPT3 (over 3%).
2. Both SDBB-NoBranch and SDBC-NoBranch yield a tighter lower-bound than TRWS, MPLP-CP-v1 and MPLP-CP-v2, which shows that the p.s.d. constraint is indeed effective.
3. Lower bounds of SDBC-NoBranch and SDBB-NoBranch converge after certain iterations. To further improve lower bounds, we need to embed the SDP bounding procedure into B&C.

Connectivity and Unary Strength We now compare the performance of SDBC, SDBB and other MAP-MRF inference methods when applied to a range of synthetic problems generated with varying parameters. We consider two factors which affect the difficulty of MAP-MRF inference problems: connectivity (*i.e.*, the average number of neighbours per node) and unary potential strength (compared to pairwise potentials). Higher connectivity and smaller unary potential strength increase the difficulty of MAP-MRF inference problems, and degrade the performance of some existing algorithms (for example, QPBO and belief propagation [116]).

In this experiment, all the synthetic MRF models have 64 variables and 10 states per variable. The pairwise energies are sampled from the Gaussian distribution $\theta_{pq}(i, j) \sim \mathcal{N}(0, 1)$, and the unary energies are sampled from $\theta_p(i) \sim \mathcal{N}(0, \omega^2)$, where the parameter ω controls the unary strength: small ω corresponds to weak unary strength, and vice versa. κ refers to the connectivity of a graph, which is the average number of linked neighbours per node.

The results are shown in Fig. 2, in which the curve of lower-bounds is the concatenation of two parts: the lower-bounds at the gradient-descent steps within the first bounding iteration and the lower-bounds for the subsequent bound-

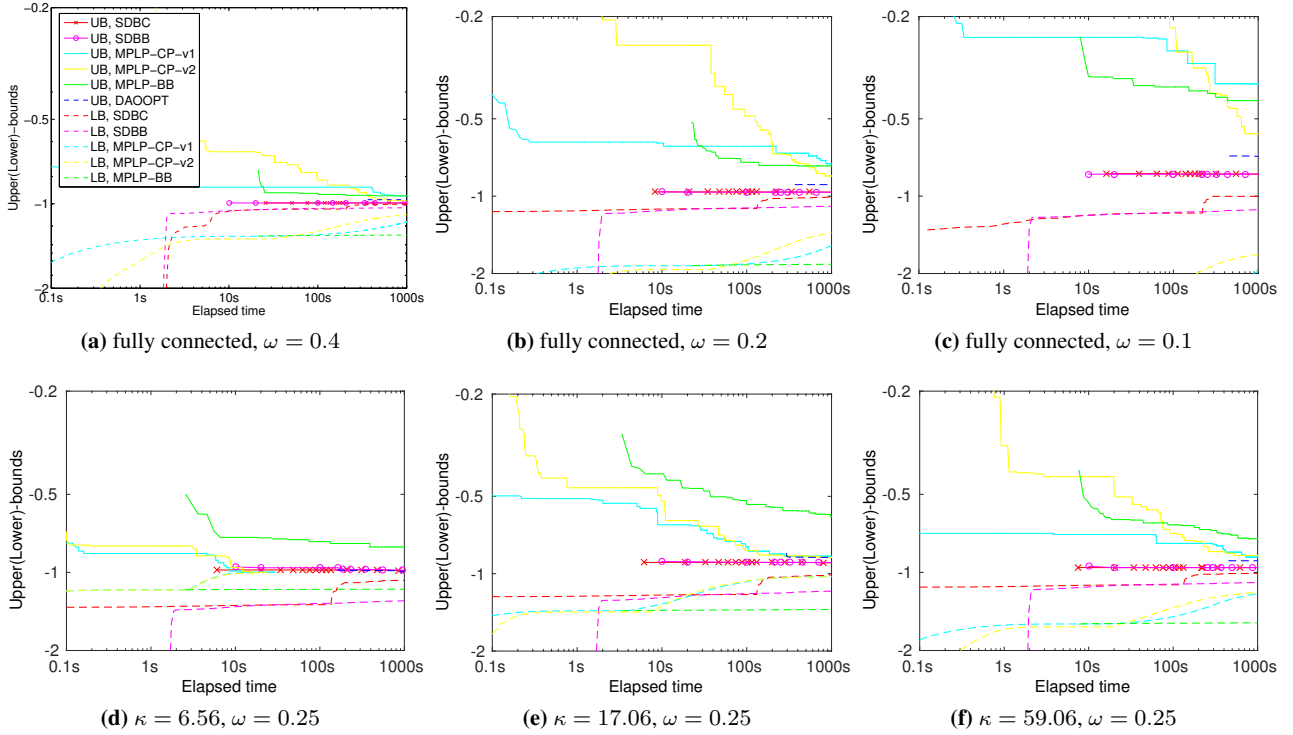


Fig. 2: Upper-/lower-bound versus running time on the synthetic MAP-MRF inference problems. All the undirected graph models in this experiment have 64 nodes and 10 states per node, and the energy functions are randomly generated. The upper/lower bound is normalized by dividing them by the best lower-bound computed within 20 minutes. In (a)-(c), the unary strength decreases from 0.4 to 0.1. When the unary strength parameter ω is large (for example, $\omega = 0.4$), all methods obtain near optimal upper-bounds (< -0.90) and the superiority of SDBC is not significant. SDBC does perform significantly better than others for smaller unary strengths (for example, $\omega = 0.1$). In (d)-(f), the connectivity κ changes from 6.56 to 59.06. For sparse graphs (for example, $\kappa = 6.56$), MPLP-CP-v1 and MPLP-CP-v2 converge quickly to the optimal and performs better than SDBC. However for dense graph models (for example, $\kappa = 59.06$), SDBC yields better upper and lower bounds.

ing iterations. In the first bounding iteration of SDBB, multiple rounding procedures are conducted at 10, 20 and 100 seconds respectively.

In Fig. 2 (a)-(c), fully-connected graphs ($\kappa = 63$) are generated and the parameter ω is set to $\{0.4, 0.2, 0.1\}$ respectively. SDBC and SDBB achieve the best upper/lower bounds in most cases. Meanwhile, SDBC yields better lower-bound than that of SDBB.

The performance of all methods becomes worse when the parameter ω decreases. However, the detrimental impact on SDBC is the smallest amongst all the evaluated algorithms. The upper-bound of SDBC drops from 0.993 ($\omega = 0.4$) to 0.820 ($\omega = 0.1$), while the upper-bound of the second best method decreases from 0.975 to 0.700. The difference on the lower-bounds is more significant: at $\omega = 0.4$, the gap between SDBC and the second best on the lower-bound is only 0.087; while the gap increases to 0.89 for $\omega = 0.1$.

In Fig. 2 (d)-(e), ω is fixed to 0.25 and the connectivity κ is set to $\{6.56, 17.06, 59.06\}$ respectively. SDBC and SDBB perform slightly worse than MPLP-CP-v1 and MPLP-CP-v2 in terms of both upper and lower bounds when $\kappa = 6.56$. For denser graphs ($\kappa \in \{17.06, 59.06\}$), however, SDBC

produces the best lower-/upper-bounds amongst all competing methods. We also compare SDBC and MPLP-CP-v1 on a larger sparse model from [117], which contains 768 variables, 8 states per variable and 2252 edges. The upper-/lower-bounds of MPLP-CP-v1 converge quickly to the global optimal solution (-565.23) using 0.8 second, as it can benefit from the sparse graph structure. While SDBC scales poorer than MPLP-CP-v1 on this type of sparse models: it obtains a non-global-optimal solution (-567.92) using around 2-hour runtime. Note that the focus of the proposed SDP approaches is on the dense problems with non-submodular terms and/or weak unary terms (See [118] for the efforts to improve the scalability of SDP algorithms for MAP inference).

Based on the experiments on this synthetic data, we may conclude that *SDBC and SDBB perform better than state-of-the-art for MRFs with high connectivity or weak unary potentials.*




Fig. 3: Image denoising. This experiment is used to validate whether SDBC converges to the globally optimal solution. The 16×16 input image (b) is generated by adding random Gaussian noise to all pixels of the ground truth image (a). The constructed graph model has one node per pixel and contains only unary and submodular pairwise potentials. Without branching, SDBC converges to the exact solution, which means the relative gap between the lower- and upper-bound is smaller than 10^{-8} . The assignment by SDBC (d) is consistent with that of the graph-cut algorithm (c). The latter is well known to provide the exact MAP solution for such binary submodular MRFs.

Kernel size	$ \mathcal{V} $	$ \mathcal{E} $	$ \mathcal{E} _{sub}/ \mathcal{E} $	$ \mathcal{V} _{redu}$
3×3	1219	13506	0.212	631.5
5×5	1219	42120	0.059	921.5
7×7	1219	82530	0.030	1094.5

Table 3: The graph sizes with respect to different convolution kernels. With the growth of kernel size, the number of edges ($|\mathcal{E}|$) increases and the portion of submodular edges ($|\mathcal{E}|_{sub}/|\mathcal{E}|$) decreases. There are more variables ($|\mathcal{V}|_{redu}$) in the reduced graph (by QPBO) for a larger kernel size.




Fig. 4: Image deconvolution. The task is to recover the binary ground truth image (first column) from the blurry and noisy gray-level input image (second column). The third column demonstrates the partial optimal results of QPBO, in which the labelled pixels are shown in black/white color and the unlabelled pixels are shown in gray color. The last two columns are the images reconstructed by MPLP-CP-v2 and SDBC, using the same computation time of 20 minutes. SDBC gives better energy value and recovery accuracy than MPLP-CP-v2.

5.2 Image Denoising and Deconvolution

The binary submodular MRF model for image denoising is generated using the UGM² code. As the energy provided by the graph-cut algorithm for such MRF models is the exact MAP value, it is used to compare with ours to validate whether the proposed B&C method converges to the exact MAP solution. As shown in Fig. 3, the result of SDBC is indeed the same as that of the graph-cut algorithm.

In this paper, we refer to image deconvolution as the task of reconstructing an image from its convolution with a *known* blurring kernel. Raj and Zabih [119] formulated the deconvolution to a MAP-MRF inference problem with unary and non-submodular pairwise energy functions. The

connectivity increases with the growth in kernel size. We use two binary images and generate the deconvolution models with different kernel sizes (3×3 , 5×5 and 7×7). QPBO is applied on the six deconvolution models to reduce the graph size, and then SDBC, MPLP-CP-v2 and DAOOPT are applied on the reduced graphs. We also report the results of QPBO and LSA-TR [107] with Euclidean or Hamming distance. LSA-TR is repeated (around 8000 times) with random initializations until the 20-minute time limit is used up. ICM is performed as a post-processing procedure for all the evaluated methods. Table 2 shows the primary evaluation results within the runtime limits of 20 minutes, 1 hour and 24 hours. DAOOPT did not give solutions within the 24 hours, so its result is omitted. Within 20 minutes, MPLP-CP-v2 converges quickly to the exact MAP solutions for the two models with respect to 3×3 kernel. But it does not converge for the rest four models with larger kernel sizes and higher connectivity. ILP fails to solve any of the six instances exactly within 24 hours. LSA-TR and QPBO provide non-optimal solutions using a much shorter time. On the other hand, SDBC gives the best upper-bounds on all the 6 instances within any of the three time limits, and is able to solve 5 instances exactly (the model with respect to “ABCD” and 7×7 is not solved exactly by SDBC).

Fig. 4 illustrates the images recovered by SDBC and MPLP-CP-v2 with respect to 7×7 kernels within 20 minutes. For image “1234”, SDBC achieves the energy value of -478.07 and the recovery accuracy of 87.4%, which are better than those given by MPLP-CP-v2 (-474.19 and 85.6%) respectively. Similarly for image “ABCD”, SDBC also performs better than MPLP-CP-v2, in terms of energy value (-416.63 vs. -410.82) and accuracy (87.5% vs. 86.0%).

We also evaluate algorithms on a larger image (Fig. 5). The binary ground truth image is convoluted by a 32×32 Gaussian kernel. The corresponding MRF problem is large: 16384 variables and 13130296 edges, so we use a parallel eigen-solver, PLASMA³. The last image shows the deconvolution result of SDBC using 5600 seconds on 16 cores. The result of MPLP-CP-v2 is demonstrated in the third image, using the same CPU hours. SDBC performs considerably better than MPLP-CP-v2.

5.3 Benchmarks

In this section, we evaluate the proposed algorithm on some MRF models in PIC2011 [108] and OpenGM [114] benchmarks. The experimental results show that our method performs better than state-of-the-arts on these models. The chosen models are difficult for many conventional methods in that these inference tasks are non-submodular, densely connected and/or with unary terms.

² <http://www.di.ens.fr/~mschmidt/Software/UGM.html>

³ <http://icl.cs.utk.edu/plasma/>

Algorithm	Upper-bound	Lower-bound	#Best-ub	#Best-lb	#Exact	Runtime
<i>Runtime $\leq 20min$</i>						
SDBC	-504.03	-505.03	6	4	0	1187.83
MPLP-CP-v2	-502.34	-514.48	2	2	2	822.46
ILP	-502.88	-568.20	1	0	0	1220.62
QPBO	-501.92	$-\infty$	0	0	0	0.76
<i>Runtime $\leq 1hr$</i>						
SDBC	-504.08	-504.45	6	4	0	3460.10
MPLP-CP-v2	-503.60	-508.40	2	2	2	2415.58
ILP	-502.91	-566.41	2	0	0	3626.06
<i>Runtime $\leq 24hr$</i>						
SDBC	-504.09	-504.09	6	6	5	41830.36
MPLP-CP-v2	-503.95	-505.37	4	2	2	57616.95
ILP	-504.07	-560.11	4	0	0	86026.42

Table 2: Image deconvolution (6 instances, 1219 variables and 2 labels for each instance). LSA-TR achieves the second best average energy value within the runtime limit of 20 minutes. MPLP-CP-v2 quickly solves two easy instances with respect to 3×3 kernel exactly, but becomes worse for models with respect to 5×5 and 7×7 kernel. ILP is not able to solve any of the six instances exactly within 24 hours, while our method achieves the best energy values for all six instances and converges to exact solutions to five instances.




Fig. 5: Image deconvolution on a 128×128 image, which is convoluted by a 32×32 Gaussian kernel. With the same running time, the result of SDBC is much better than MPLP-CP-v2.

Algorithm	Upper-bound	Lower-bound	#Best-ub	#Best-lb	#Exact
<i>Runtime = 20min</i>					
SDBC	-19302.83	-19910.31	34	37	0
MPLP-CP-v1	-18595.15	-24749.11	10	0	0
MPLP-CP-v2	-18905.71	-24624.95	13	0	0
ILP	-18247.73	-542806.60	8	0	0
DAOOPT	-19139.70	$-\infty$	17	0	0
MPLP-BB	-17838.86	-36904.82	11	0	0
<i>Runtime = 1hr</i>					
SDBC	-19303.06	-19650.62	35	35	0
MPLP-CP-v1	-18985.04	-22312.83	17	1	0
MPLP-CP-v2	-19231.27	-22260.29	19	1	0
ILP	-19049.54	-34838.62	15	0	0
DAOOPT	-19139.70	$-\infty$	17	0	0
MPLP-BB	-18227.13	-36804.25	10	0	0
<i>Runtime = 24hr</i>					
SDBC	-19313.25	-19553.50	36	34	3
MPLP-CP-v2	-19270.70	-21138.12	22	3	0
ILP	-19221.27	-33888.65	22	0	0
DAOOPT	-19139.70	$-\infty$	16	0	0

Table 4: PIC2011-ObjectDetection (37 instances, 60 variables and 10-20 labels for each instance). SDBC achieves the best upper-bounds and lower-bounds on most of instances within 20 minutes, 1 hour and 24 hours. Exact solutions to 3 instances are obtained by our algorithm.

5.3.1 PIC2011-ObjectDetection

There are 37 models in the “Object-Detection” sub-category of the PIC2011 challenge, which are fully-connected, lack-

ing in unary terms and have 60 nodes and 10/15/20 states per node. SDBC is compared with MPLP-CP-v1, MPLP-CP-v2, ILP, DAOOPT and MPLP-BB, among which DAOOPT is the winner of the PIC2011 challenge.

The primary results are shown in Table 4. The upper-bounds and lower-bounds of different methods are compared under three running-time limits: 20 minutes, 1 hour and 24 hour. Our method achieves the best upper- and lower-bounds for most of the 37 models under all time limits. DAOOPT is the second best method within 20 minutes, while MPLP-CP-v2 becomes better than DAOOPT within 1 hour. ILP performs poorly within the runtime limit of 20 minutes, while it becomes the second best in terms of the number of best upper-bounds (22, the same as MPLP-CP-v2).

In particular, all the evaluated methods do not converge to global optimal solutions on these models, except that SDBC gives exact solutions to three of the 37 models.

5.3.2 OpenGM-ChineseChar

The graphical model for Chinese character inpainting in the OpenGM benchmarks contains non-submodular terms, such that graph cuts based methods cannot give exact solutions in general. It is shown in [1] that applying combinatorial algorithms (like MCBC [111] and ILP [109]) on the models reduced by QPBO achieves the best performance. Likewise, we also apply SDBC on the models reduced by QPBO and achieve comparable results. Within one hour, our method solves 44 out of the 100 instances exactly, which is only worse than the specialized solver, MCBC. Note that MCBC utilizes more complicated tightening constraints and rounding approach than ours. On the other hand, our method is better than all the other methods (including MCBC) in terms of the number of best upper-bounds (84) and lower-bounds (85). In particular, SDBC achieves better lower-bounds on those instances the LP methods (MCBC and ILP) are loose, which demonstrate the potential of SDBC in solving these instances exactly. To validate this potential, we run SDBC for 4 hours and find that it gives global optimal solutions of 80 instances.

5.3.3 OpenGM-ModularityClustering

The graphical models in this subclass is to find a clustering of a network maximizing the modularity. Although the six instances have small graph sizes (34 to 115 variables), they are difficult due to the absence of unary terms and fully-connected graph structure.

It can be seen in [1] that the LP relaxation method with cycle and odd-wheel constraints (MCR-CCFDB-OWC) [112] and the LP/ILP method (MCI-CCFDB-CCIFD) [113] perform the best on modularity clustering, which solve five of the six instances exactly. However, they do not converge on

the largest instance within one hour. Kernighan-Lin (KL) algorithm [110], a specialized efficient heuristic for clustering networks, offers a better solution than MCI and MCR on this instance. As MCR and MCI, our method also solves these five instances within one hour *without branching* (in other words, using only one bounding procedure). Moreover, SDBC provides the best solution (in terms of upper-bound) to the largest instance among all the evaluated approaches.

6 Conclusion

We have presented an efficient Branch-and-Cut method for MAP-MRF inference problems by taking the advantage of an efficient and tight SDP bounding procedure. The main contribution of the proposed method include: 1) A variety of linear tightening constraints have been incorporated in SDP relaxation to MAP problems, by using an efficient SDP solver and cutting-plane. 2) Several techniques have also been employed to make the cutting-plane and branch-and-bound more efficient, including model reduction, warm start and dropping inactive constraints. Experiments demonstrate that the proposed method performs very well, particularly for problems with high connectivity or weak unary potentials. These types of problems are difficult to solve by almost all other existing methods. Furthermore, our method usually provide good approximate solutions within the first several bounding procedures. In the experiments, we have compared the proposed approach with state-of-the-art methods. The results demonstrate the superior performance of our approach on the evaluated problems.

7 Appendix

7.1 Relationship between the Standard SDP Relaxation (11) and the Simplified Dual (21)

The Lagrangian dual of (11) can be expressed in the following general form:

$$\min_{\mathbf{u}} \quad \mathbf{u}^\top \mathbf{b} \quad (25a)$$

$$\text{s.t.} \quad \mathbf{Z} = \mathbf{A} + \sum_{i=1}^m u_i \mathbf{B}_i \succcurlyeq \mathbf{0}, \quad (25b)$$

$$u_i \geq 0, \forall i \in \mathcal{I}_{in}. \quad (25c)$$

The p.s.d. constraint (25b) can be replaced by a *penalty function*, which is considered as a measure of violation of this constraint. In our case, the penalty function is defined as $p(\mathbf{u}) = \|\min(\mathbf{0}, \boldsymbol{\lambda})\|_2^2 = \|\Pi_{\mathcal{S}_+^{nh+1}}(\mathbf{C}(\mathbf{u}))\|_F^2$, where $\boldsymbol{\lambda}$ is the vector of eigenvalues of \mathbf{Z} . We can find that if $p(\mathbf{u}) = 0$, then $\mathbf{Z} \succcurlyeq \mathbf{0}$. Now the problem (25) can be transformed to

$$\min_{\mathbf{u}} \quad \mathbf{u}^\top \mathbf{b} + \frac{\gamma}{2} \|\Pi_{\mathcal{S}_+^{nh+1}}(\mathbf{C}(\mathbf{u}))\|_F^2 \quad (26a)$$

$$\text{s.t.} \quad u_i \geq 0, \forall i \in \mathcal{I}_{in}, \quad (26b)$$

Algorithm	Upper-bound	Lower-bound	#Best-ub	#Best-lb	#Exact	Runtime (s)
SDBC	-49546.3380	-49558.9258	84	85	44	2332.18
MCBC	-49550.0972	-49612.3817	72	57	56	2053.89
ILP	-49547.4144	-50061.1506	59	0	0	3553.71
LSA-TR(euc.)	-49548.0996	$-\infty$	28	0	0	0.05
LSA-TR(ham.)	-49536.7646	$-\infty$	1	0	0	0.06
QPBO	-49501.9531	-50119.3835	0	0	0	0.17

Table 5: OpenGM-ChineseChar (100 instances, 8000 variables and 2 labels for each instance). For fair comparison, all the combinatorial methods (SDBC, MCBC and ILP) are evaluated on graphs reduced by QPBO. The results of all algorithms except SDBC are obtained from [1]. Within 1 hour, SDBC solves 44 instances exactly and achieves the best upper-bounds and lower-bounds on over 80% instances.

Algorithm	Upper-bound	Lower-bound	#Best-ub	#Best-lb	#Exact	Runtime (s)
SDBC	-0.4913	-0.4939	6	6	5	1500.26
KL	-0.4860	$-\infty$	2	0	0	0.01
MCR-CCFDB-OWC	-0.4400	-0.5021	5	5	5	601.38
MCI-CCFDB-CCIFD	-0.4652	-0.4962	5	5	5	602.75
MCI-CCI	-0.4312	-0.5158	4	4	4	1207.07
MCI-CCIFD	-0.4399	-0.5176	4	4	4	1204.03

Table 6: OpenGM-ModularityClustering (6 instances, 34-115 variables and 34-115 labels for each instance). As the top-performing algorithms (MCR-CCFDB-OWC and MCI-CCFDB-CCIFD) in [1], SDBC solves 5 instances exactly. On the remaining largest instance, our method achieves the best lower-bound and upper-bound over all compared methods within 1 hour.

Dataset	Solution	#Instances	#Bounding	#Pruned	$ \mathcal{Q} $	Runtime
Deconvolution	Exact	5	10.80	5.90	0.00	33332
(6 instances, runtime ≤ 24 hr)	Inexact	1	19.00	6.00	8.00	84324
Object-Detection	Exact	3	75.67	38.33	0.00	33625
(37 instances, runtime ≤ 24 hr)	Inexact	34	74.68	19.50	36.68	85631
ChineseChar	Exact	80	3.42	2.21	0.00	4347
(100 instances, runtime ≤ 4 hr)	Inexact	20	3.45	1.00	2.45	12392
Modularity-Clustering	Exact	5	1.00	1.00	0.00	1063
(6 instances, runtime ≤ 1 hr)	Inexact	1	1.00	0.00	2.00	3686

Table 7: The average number of bounding iterations, pruned subproblems and unsolved subproblems in the queue \mathcal{Q} for SDBC. The results for instances solved exactly and inexactly are summarized individually. SDBC uses 1-10 bounding iterations to reach the exact solutions to most of instances on Deconvolution, ChineseChar and Modularity-Clustering. For Object-Detection, SDBC is able to solve 3 instances exactly with less than 100 bounding evaluations.

where $\gamma > 0$ serves as a penalty parameter. With the increase of γ , the solution to (26) converges to that of (25). It is clear that (26) is equivalent to (21).

7.2 Proof of Propositions 1 and 2

Firstly, It is known [49, 120] that the set of p.s.d. matrices with fixed trace $\Theta_\eta := \{\mathbf{X} \succeq \mathbf{0} | \text{trace}(\mathbf{X}) = \eta\}$, $\forall \eta > 0$ has the following property:

Theorem 2 (*The spherical constraint*). $\forall \eta > 0, \forall \mathbf{X} \in \Theta_\eta$, we have $\|\mathbf{X}\|_F \leq \eta$, and $\|\mathbf{X}\|_F = \eta$ if and only if $\text{rank}(\mathbf{X}) = 1$.

It is also shown in [49] that the problem (21) is the Lagrangian dual of the following problem:

$$\min_{\mathbf{y}, \mathbf{Y}} E(\mathbf{y}, \mathbf{Y}) + g_\gamma(\mathbf{y}, \mathbf{Y}) \quad (27a)$$

$$\text{s.t. (12), (13), (14), (15), (16), (17), (18), (19),} \quad (27b)$$

$$\Omega(\mathbf{y}, \mathbf{Y}) \succeq \mathbf{0}, \quad (27c)$$

where $g_\gamma(\mathbf{y}, \mathbf{Y}) = \frac{1}{2\gamma} (\|\Omega(\mathbf{y}, \mathbf{Y})\|_F^2 - (n+1)^2)$.

Proof of Proposition 1: (i) $\forall \mathcal{D}_1 \subseteq \mathcal{D}_2 \subseteq \mathcal{Z}^n$, $\exists \mathcal{F}_{in}, \mathcal{F}_{eq} \in \{(p, i)\}_{p \in \mathcal{V}, i \in \mathcal{Z}}$ such that $\mathcal{D}_1 = \{\mathbf{x} \in \mathcal{D}_2 \mid x_p \neq i, \forall (p, i) \in \mathcal{F}_{in}; x_p = i, \forall (p, i) \in \mathcal{F}_{eq}\}$. Consequently, the difference between the SDCut primal formulation (27) with respect to \mathcal{D}_1 and \mathcal{D}_2 is that the one with respect to \mathcal{D}_1 contains the following additional linear constraints:

$$\begin{cases} y_{p,i} = 0, Y_{pi,qj} = Y_{qj,pi} = 0, & \forall (p, i) \in \mathcal{F}_{in}, \\ y_{p,i} = 1, Y_{pi,qj} = Y_{qj,pi} = y_{q,j}, & \forall (p, i) \in \mathcal{F}_{eq}. \end{cases} \quad (28)$$

Because of the strong duality, we know that $d_\gamma^*(\mathcal{D})$ equals to the optimal value of the corresponding primal problem (27). Then we have $d_\gamma^*(\mathcal{D}_1) \geq d_\gamma^*(\mathcal{D}_2)$, as the primal problem (27) with respect to \mathcal{D}_1 has more constraints than that with respect to \mathcal{D}_2 .

(ii) This proof is simple. As $|\mathcal{D}| = 1$, there is only one point $\hat{\mathbf{x}}$ in the set \mathcal{D} and $E(\hat{\mathbf{x}}) = \min_{\mathbf{x} \in \mathcal{D}} E(\mathbf{x})$. Then the feasible set of (27) also contains a single point $\{\hat{\mathbf{y}}, \hat{\mathbf{Y}}\}$ corresponding to $\hat{\mathbf{x}}$ by applying constraints as (28). Because $\|\Omega(\hat{\mathbf{y}}, \hat{\mathbf{Y}})\|_F^2 = (n+1)^2$, we have $d_\gamma^*(\mathcal{D}) = E(\hat{\mathbf{y}}, \hat{\mathbf{Y}}) = \min_{\mathbf{x} \in \mathcal{D}} E(\mathbf{x})$.

Proof of Proposition 2: $\{\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*\}$ is the optimal solution of (27) based on the strong duality, and $d_\gamma(\mathbf{u}_\gamma^*)$ is the corresponding optimal objective value. Consider the following problem

$$\min_{\mathbf{y}, \mathbf{Y}} E(\mathbf{y}, \mathbf{Y}) + g_\gamma(\mathbf{y}, \mathbf{Y}) \quad (29a)$$

$$\text{s.t. } \Omega(\mathbf{y}, \mathbf{Y}) \succeq \mathbf{0}, \text{rank}(\Omega(\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*)) = 1, \quad (29b)$$

$$(12), (13), (14), (15), (16), (17), (18), (19), \quad (29c)$$

which adds a rank-1 constraint to the problem (27). Then $d_\gamma(\mathbf{u}_\gamma^*)$ and $\{\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*\}$ are also optimal for the above problem. Note that the constraints (12), (13), $\Omega(\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*) \succeq \mathbf{0}$ and $\text{rank}(\Omega(\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*)) = 1$, force $\{\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*\}$ to be a vertex of $\mathcal{M}(\mathcal{G}, \mathcal{Z})$. So the feasible set of (29) is $\mathcal{M}(\mathcal{G}, \mathcal{Z})$. On the other hand, $g_\gamma(\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*) = 0$ at $\text{rank}(\Omega(\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*)) = 1$ (Theorem 2), so the objective function of (29) is $E(\mathbf{y}, \mathbf{Y})$. In summary, the problem (29) is equivalent to the MAP problem $\min_{\mathbf{y}, \mathbf{Y} \in \mathcal{M}(\mathcal{G}, \mathcal{Z})} E(\mathbf{y}, \mathbf{Y})$. Then we have that $\mathbf{y}_\gamma^*, \mathbf{Y}_\gamma^*$ yield the exact MAP solution and $d_\gamma(\mathbf{u}_\gamma^*)$ is the minimum energy. The value of γ does not affect the above results.

References

1. J. H. Kappes, B. Andres, F. A. Hamprecht, C. Schnörr, S. Nowozin, D. Batra, S. Kim, B. X. Kausler, T. Kröger, J. Lellmann, N. Komodakis, B. Savchynskyy, and C. Rother, “A comparative study of modern inference techniques for structured discrete energy minimization problems,” *Int. J. Comp. Vis.*, 2015.
2. R. Szeliski, R. Zabih, D. Scharstein, O. Veksler, V. Kolmogorov, A. Agarwala, M. Tappen, and C. Rother, “A comparative study of energy minimization methods for markov random fields with smoothness-based priors,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 6, pp. 1068–1080, 2008.
3. M. P. Kumar, V. Kolmogorov, and P. H. S. Torr, “An analysis of convex relaxations for MAP estimation of discrete MRFs,” *J. Mach. Learn. Res.*, vol. 10, pp. 71–106, Jun 2009.
4. V. Kolmogorov and C. Rother, “Comparison of energy minimization algorithms for highly connected graphs,” in *Proc. Eur. Conf. Comp. Vis.*, 2006, pp. 1–15.
5. S. D. Givry, B. Hurley, D. Allouche, G. Katsirelos, B. O’Sullivan, and T. Schiex, “An experimental evaluation of CP/AI/OR solvers for optimization in graphical models,” in *Congrès ROADEF’2014, Bordeaux, FRA*, 2014.
6. Y. Boykov, O. Veksler, and R. Zabih, “Fast approximate energy minimization via graph cuts,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 11, pp. 1222–1239, 2001.
7. V. Kolmogorov and R. Zabih, “What energy functions can be minimized via graph cuts?” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 2, pp. 147–159, 2004.
8. V. Kolmogorov and C. Rother, “Minimizing nonsubmodular functions with graph cuts—a review,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 7, pp. 1274–1279, 2007.
9. C. Rother, V. Kolmogorov, V. Lempitsky, and M. Szummer, “Optimizing binary MRFs via extended roof duality,” in *Proc. IEEE Conf. Comp. Vis. Pattern Recogn.*, 2007, pp. 1–8.
10. J. Pearl, *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.
11. Y. Weiss and W. T. Freeman, “On the optimality of solutions of the max-product belief-propagation algorithm in arbitrary graphs,” *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 736–744, 2001.
12. J. Sun, H.-Y. Shum, and N.-N. Zheng, “Stereo matching using belief propagation,” in *Proc. Eur. Conf. Comp. Vis.*, 2002, pp. 510–524.
13. P. F. Felzenszwalb and D. P. Huttenlocher, “Efficient belief propagation for early vision,” *Int. J. Comp. Vis.*, vol. 70, no. 1, pp. 41–54, 2006.
14. S. M. Aji, G. B. Horn, and R. J. McEliece, “On the convergence of iterative decoding on graphs with a single cycle,” in *Proc. ISIT*, 1998.
15. Y. Weiss, “Correctness of local probability propagation in graphical models with loops,” *Neural Comp.*, vol. 12, no. 1, 2000.
16. M. Bayati, D. Shah, and M. Sharma, “Maximum weight matching via max-product belief propagation,” in *Proc. IEEE Intl. Symp. on Information Theory*. IEEE, 2005, pp. 1763–1767.
17. M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky, “MAP estimation via agreement on trees: message-passing and linear programming,” *IEEE Trans. Inf. Theory*, vol. 51, no. 11, pp. 3697–3717, 2005.
18. V. Kolmogorov, “Convergent tree-reweighted message passing for energy minimization,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 10, pp. 1568–1583, 2006.
19. V. Kolmogorov and M. J. Wainwright, “On the optimality of tree-reweighted max-product message-passing,” in *Proc. Uncertainty in Artificial Intell.*, 2005.
20. M. Shlezinger, “Syntactic analysis of two-dimensional visual signals in noisy conditions,” *Kibernetika*, vol. 4, pp. 113–130, 1976.
21. 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, 2007.
22. M. J. Wainwright and M. I. Jordan, “Graphical models, exponential families, and variational inference,” *Found.*

- dations and Trends® in Machine Learning*, vol. 1, no. 1-2, pp. 1–305, 2008.
23. M. P. Kumar, P. H. Torr, and A. Zisserman, “Solving Markov random fields using second order cone programming relaxations,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, vol. 1, 2006, pp. 1045–1052.
 24. P. Ravikumar and J. Lafferty, “Quadratic programming relaxations for metric labeling and Markov random field MAP estimation,” in *Proc. Int. Conf. Mach. Learn.*, 2006, pp. 737–744.
 25. A. Globerson and T. S. Jaakkola, “Fixing max-product: Convergent message passing algorithms for MAP LP-relaxations,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2007.
 26. 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, 2011.
 27. J. H. Kappes, S. Schmidt, and C. Schnörr, “MRF inference by k-fan decomposition and tight lagrangian relaxation,” in *Proc. Eur. Conf. Comp. Vis.*, 2010, pp. 735–747.
 28. V. Jojic, S. Gould, and D. Koller, “Accelerated dual decomposition for MAP inference,” in *Proc. Int. Conf. Mach. Learn.*, 2010, pp. 503–510.
 29. J. H. Kappes, B. Savchynskyy, and C. Schnörr, “A bundle approach to efficient MAP-inference by lagrangian relaxation,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2012, pp. 1688–1695.
 30. P. Ravikumar, A. Agarwal, and M. J. Wainwright, “Message-passing for graph-structured linear programs: Proximal methods and rounding schemes,” *J. Mach. Learn. Res.*, vol. 11, pp. 1043–1080, 2010.
 31. A. F. Martins, N. A. Smith, E. P. Xing, P. M. Aguiar, and M. A. Figueiredo, “Augmenting dual decomposition for MAP inference,” in *Proc. Int. Workshop on Optimization for Machine Learning*, 2011.
 32. A. F. Martins, M. A. Figueiredo, P. M. Aguiar, N. A. Smith, and E. P. Xing, “An augmented Lagrangian approach to constrained MAP inference,” in *Proc. Int. Conf. Mach. Learn.*, 2011, pp. 169–176.
 33. O. Meshi and A. Globerson, “An alternating direction method for dual MAP LP relaxation,” in *Machine Learning and Knowledge Discovery in Databases*. Springer, 2011, pp. 470–483.
 34. 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, 2010.
 35. J. K. Johnson, D. M. Malioutov, and A. S. Willsky, “Lagrangian relaxation for MAP estimation in graphical models,” *45th Ann. Allerton Conf. on Comm., Control and Comp.*, 2007.
 36. B. Savchynskyy, S. Schmidt, J. Kappes, and C. Schnörr, “Efficient MRF energy minimization via adaptive diminishing smoothing,” *Proc. Uncertainty in Artificial Intell.*, 2012.
 37. N. Komodakis and N. Paragios, “Beyond loose LP-relaxations: Optimizing MRFs by repairing cycles,” in *Proc. Eur. Conf. Comp. Vis.*, 2008, pp. 806–820.
 38. T. Werner, “High-arity interactions, polyhedral relaxations, and cutting plane algorithm for soft constraint optimisation (MAP-MRF),” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2008, pp. 1–8.
 39. D. Sontag, T. Meltzer, A. Globerson, T. S. Jaakkola, and Y. Weiss, “Tightening LP relaxations for MAP using message passing,” in *Proc. Uncertainty in Artificial Intell.*, 2008.
 40. D. Batra, S. Nowozin, and P. Kohli, “Tighter relaxations for MAP-MRF inference: A local primal-dual gap based separation algorithm,” in *Proc. Int. Conf. Artificial Intell. & Statistics*, 2011, pp. 146–154.
 41. J. K. Johnson, “Convex relaxation methods for graphical models: Lagrangian and maximum entropy approaches,” Ph.D. dissertation, Massachusetts Institute of Technology, 2008.
 42. D. Sontag, D. K. Choe, and Y. Li, “Efficiently searching for frustrated cycles in MAP inference,” in *Proc. Uncertainty in Artificial Intell.*, 2012.
 43. P. H. S. Torr, “Solving Markov random fields using semi definite programming,” in *Proc. Int. Conf. Artificial Intell. & Statistics*, 2003.
 44. M. J. Jordan and M. I. Wainwright, “Semidefinite relaxations for approximate inference on graphs with cycles,” in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 16, 2003, pp. 369–376.
 45. C. Olsson, A. P. Eriksson, and F. Kahl, “Solving large scale binary quadratic problems: Spectral methods vs. semidefinite programming,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2007, pp. 1–8.
 46. C. Schellewald and C. Schnörr, “Probabilistic subgraph matching based on convex relaxation,” in *Workshop of IEEE Conf. Comp. Vis. Patt. Recogn.*, 2005, pp. 171–186.
 47. A. Joulin, F. Bach, and J. Ponce, “Discriminative clustering for image co-segmentation,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2010.
 48. J. Peng, T. Hazan, N. Srebro, and J. Xu, “Approximate inference by intersecting semidefinite bound and local polytope,” in *Proc. Int. Conf. Artificial Intell. & Statistics*, 2012, pp. 868–876.
 49. P. Wang, C. Shen, and A. van den Hengel, “A fast semidefinite approach to solving binary quadratic problems,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2013.

50. Q. Huang, Y. Chen, and L. Guibas, “Scalable semidefinite relaxation for maximum a posterior estimation,” in *Proc. Int. Conf. Mach. Learn.*, 2014.
51. R. Frostig, S. Wang, P. S. Liang, and C. D. Manning, “Simple MAP inference via low-rank relaxations,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 3077–3085.
52. D. Sontag and T. S. Jaakkola, “New outer bounds on the marginal polytope,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2007.
53. M. X. Goemans and D. P. Williamson, “Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming,” *J. ACM*, vol. 42, pp. 1115–1145, 1995.
54. F. Alizadeh, J.-P. A. Haeberly, and M. L. Overton, “Primal-dual interior-point methods for semidefinite programming: convergence rates, stability and numerical results,” *SIAM J. Optim.*, vol. 8, no. 3, pp. 746–768, 1998.
55. E. D. Andersen, C. Roos, and T. Terlaky, “On implementing a primal-dual interior-point method for conic quadratic optimization,” *Math. Program.*, vol. 95, no. 2, pp. 249–277, 2003.
56. Y. E. Nesterov and M. J. Todd, “Primal-dual interior-point methods for self-scaled cones,” *SIAM J. Optim.*, vol. 8, no. 2, pp. 324–364, 1998.
57. Y. Ye, M. J. Todd, and S. Mizuno, “An $O(\sqrt{nL})$ -iteration homogeneous and self-dual linear programming algorithm,” *Mathematics of Operations Research*, vol. 19, no. 1, pp. 53–67, 1994.
58. C. Helmberg, “An interior point method for semidefinite programming and max-cut bounds,” Ph.D. dissertation, Department of Mathematics, Graz University of Technology, 1994.
59. S. Burer and R. D. Monteiro, “A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization,” *Math. Program.*, vol. 95, no. 2, pp. 329–357, 2003.
60. S. Arora, E. Hazan, and S. Kale, “The multiplicative weights update method: a meta-algorithm and applications,” *Theory of Computing*, vol. 8, no. 1, pp. 121–164, 2012.
61. —, “Fast algorithms for approximate semidefinite programming using the multiplicative weights update method,” in *Proc. Annual IEEE Symposium on Foundations of Computer Science*, 2005, pp. 339–348.
62. S. Arora and S. Kale, “A combinatorial, primal-dual approach to semidefinite programs,” in *Proc. Annual ACM symposium on Theory of computing*, 2007, pp. 227–236.
63. E. Hazan, “Sparse approximate solutions to semidefinite programs,” in *LATIN 2008: Theoretical Informatics*, 2008, pp. 306–316.
64. D. Garber and E. Hazan, “Approximating semidefinite programs in sublinear time,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2011, pp. 1080–1088.
65. S. Laue, “A hybrid algorithm for convex semidefinite optimization,” in *Proc. Int. Conf. Mach. Learn.*, 2012, pp. 177–184.
66. J. Malick, J. Povh, F. Rendl, and A. Wiegele, “Regularization methods for semidefinite programming,” *SIAM J. Optim.*, vol. 20, no. 1, pp. 336–356, 2009.
67. X.-Y. Zhao, D. Sun, and K.-C. Toh, “A Newton-CG augmented Lagrangian method for semidefinite programming,” *SIAM J. Optim.*, vol. 20, no. 4, pp. 1737–1765, 2010.
68. Z. Wen, D. Goldfarb, and W. Yin, “Alternating direction augmented Lagrangian methods for semidefinite programming,” *Mathematical Programming Computation*, vol. 2, no. 3–4, pp. 203–230, 2010.
69. R. T. Rockafellar, “A dual approach to solving nonlinear programming problems by unconstrained optimization,” *Math. Program.*, vol. 5, no. 1, pp. 354–373, 1973.
70. C. Shen, J. Kim, and L. Wang, “A scalable dual approach to semidefinite metric learning,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2011, pp. 2601–2608.
71. M. Sun, M. Telaprolu, H. Lee, and S. Savarese, “Efficient and exact MAP inference using branch and bound,” in *Proc. Int. Conf. Artificial Intell. & Statistics*, 2012.
72. L. Otten and R. Dechter, “Anytime AND/OR depth-first search for combinatorial optimization,” *AI Communications*, vol. 25, no. 3, pp. 211–227, 2012.
73. N. Rislock, J. Malick, and F. Roupin, “Improved semidefinite bounding procedure for solving max-cut problems to optimality,” *Math. Program.*, 2012.
74. M. Armbruster, M. Fügenschuh, C. Helmberg, and A. Martin, “LP and SDP branch-and-cut algorithms for the minimum graph bisection problem: a computational comparison,” *Mathematical Programming Computation*, vol. 4, no. 3, pp. 275–306, 2012.
75. S. Burer and D. Vandenbussche, “A finite branch-and-bound algorithm for nonconvex quadratic programming via semidefinite relaxations,” *Math. Program.*, vol. 113, no. 2, pp. 259–282, 2008.
76. T. P. Dinh, N. N. Canh, and H. A. Le Thi, “An efficient combined DCA and B&B using DC/SDP relaxation for globally solving binary quadratic programs,” *Journal of Global Optimization*, vol. 48, no. 4, pp. 595–632, 2010.
77. S. Mars and L. Schewe, “SDP-package for SCIP,” TU Darmstadt, Tech. Rep., 2012.
78. C. Helmberg and R. Weismantel, “Cutting plane algorithms for semidefinite relaxations,” *Fields Institute Communications*, vol. 18, pp. 197–213, 1998.

79. C. Helmberg, F. Rendl, and R. Weismantel, "A semidefinite programming approach to the quadratic knapsack problem," *Journal of combinatorial optimization*, vol. 4, no. 2, pp. 197–215, 2000.
80. —, "Quadratic knapsack relaxations using cutting planes," in *Proceedings of the 5th International IPCO Conference on Integer Programming and Combinatorial Optimization*, 1996, pp. 175–189.
81. C. Helmberg and F. Rendl, "Solving quadratic (0, 1)-problems by semidefinite programs and cutting planes," *Mathematical Programming*, vol. 82, no. 3, pp. 291–315, 1998.
82. C. Helmberg, S. Poljak, F. Rendl, and H. Wolkowicz, "Combining semidefinite and polyhedral relaxations for integer programs," in *Proceedings of the 4th International IPCO Conference on Integer Programming and Combinatorial Optimization*, 1995, pp. 124–134.
83. J. Besag, "On the statistical analysis of dirty pictures," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 259–302, 1986.
84. R. Horst, P. M. Pardalos, and N. Van Thoai, *Introduction to global optimization*. Springer, 2000.
85. E. M. Hendrix, G. Boglárka et al., *Introduction to non-linear and global optimization*. Springer, 2010.
86. T. Achterberg, T. Koch, and A. Martin, "Branching rules revisited," *Operations Research Letters*, vol. 33, no. 1, pp. 42–54, 2005.
87. A. Land and S. Powell, "Computer codes for problems of integer programming," *Ann. Discrete Math.*, vol. 5, pp. 221–269, 1979.
88. J. T. Linderoth and M. W. Savelsbergh, "A computational study of search strategies for mixed integer programming," *INFORMS Journal on Computing*, vol. 11, no. 2, pp. 173–187, 1999.
89. G. Mitra, "Investigation of some branch and bound strategies for the solution of mixed integer linear programs," *Mathematical Programming*, vol. 4, no. 1, pp. 155–170, 1973.
90. F. Barahona and A. R. Mahjoub, "On the cut polytope," *Mathematical programming*, vol. 36, no. 2, pp. 157–173, 1986.
91. M. Deza, M. Grötschel, and M. Laurent, "Clique-web facets for multicut polytopes," *Mathematics of Operations Research*, vol. 17, no. 4, pp. 981–1000, 1992.
92. S. Chopra and M. R. Rao, "The partition problem," *Mathematical Programming*, vol. 59, no. 1-3, pp. 87–115, 1993.
93. M. Deza and M. Laurent, "Geometry of cuts and metrics, volume 15 of algorithms and combinatorics," 1997.
94. C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, "L-BFGS-B: Fortran subroutines for large-scale bound constrained optimization," *ACM T. Math. Softw.*, 1997.
95. I. Kovtun, "Sufficient condition for partial optimality for (max,+)-labeling problems and its usage," *Control systems and Computers*, vol. 2, 2011.
96. P. Swoboda, B. Savchynskyy, J. Kappes, and C. Schnörr, "Partial optimality via iterative pruning for the potts model," in *SSVM*, 2013.
97. P. Kohli, A. Shekhovtsov, C. Rother, V. Kolmogorov, and P. Torr, "On partial optimality in multi-label mrfs," in *Proc. Int. Conf. Mach. Learn.*, 2008, pp. 480–487.
98. T. Windheuser, H. Ishikawa, and D. Cremers, "Generalized roof duality for multi-label optimization: optimal lower bounds and persistency," in *Proc. Eur. Conf. Comp. Vis.*, 2012, pp. 400–413.
99. 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. Comp. Vis. Patt. Recogn.*, 2014, pp. 1170–1177.
100. A. Shekhovtsov, "Maximum persistency in energy minimization," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2014, pp. 1162–1169.
101. A. Shekhovtsov, P. Swoboda, B. Savchynskyy, I. used by Alahari et al., "Maximum persistency via iterative relaxed inference with graphical models," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2015, pp. 521–529.
102. K. Alahari, P. Kohli, and P. H. Torr, "Reduce, reuse & recycle: Efficiently solving multi-label MRFs," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2008, pp. 1–8.
103. 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. Comp. Vis. Patt. Recogn.*, 2013, pp. 1752–1758.
104. D. Topkis, "A cutting-plane algorithm with linear and geometric rates of convergence," *Journal of Optimization Theory and Applications*, vol. 36, no. 1, pp. 1–22, 1982.
105. T. Joachims, T. Finley, and C.-N. J. Yu, "Cutting-plane training of structural svms," *Machine Learning*, vol. 77, no. 1, pp. 27–59, 2009.
106. R. Horst and H. Tuy, *Global optimization: Deterministic approaches*. Springer, 2013.
107. L. Gorelick, Y. Boykov, O. Veksler, I. B. Ayed, and A. Delong, "Local submodular approximations for binary pairwise energies," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2014.
108. G. Elidan and A. Globerson, "The probabilistic inference challenge (PIC2011)," <http://www.cs.huji.ac.il/project/PASCAL/>, 2011.
109. IBM, "ILOG CPLEX optimizer," <http://www-01.ibm.com/software/integration/optimization/cplex->

- optimizer/, 2015.
110. B. Kernighan and S. Lin, “An efficient heuristic procedure for partitioning graphs,” *The Bell Systems Technical Journal*, vol. 49, no. 2, pp. 291–307, 1970.
 111. T. Bonato, M. Jünger, G. Reinelt, and G. Rinaldi, “Lifting and separation procedures for the cut polytope,” *Mathematical Programming*, vol. 146, no. 1-2, pp. 351–378, 2014.
 112. J. H. Kappes, M. Speth, B. Andres, G. Reinelt, and C. Schnörr, “Globally optimal image partitioning by multicuts,” in *Energy Minimization Methods in Computer Vision and Pattern Recognition*, 2011, pp. 31–44.
 113. J. H. Kappes, M. Speth, G. Reinelt, and C. Schnörr, “Higher-order segmentation via multicuts,” *arXiv preprint arXiv:1305.6387*, 2013.
 114. B. Andres, T. Beier, and J. H. Kappes, “OpenGM2,” <http://hci.iwr.uni-heidelberg.de/opengm2/>, 2014.
 115. R. H. Tütüncü, K. C. Toh, and M. J. Todd, “Solving semidefinite-quadratic-linear programs using SDPT3,” *Math. Program.*, vol. 95, no. 2, pp. 189–217, 2003.
 116. W. Kim and K. M. Lee, “A hybrid approach for MRF optimization problems: Combination of stochastic sampling and deterministic algorithms,” *Comp. Vis. Image Understanding*, vol. 115, no. 12, pp. 1623–1637, 2011.
 117. F. Liu, G. Lin, and C. Shen, “CRF learning with CNN features for image segmentation,” *Pattern Recognition*, 2015.
 118. P. Wang, C. Shen, and A. V. D. Hengel, “Efficient SDP inference for fully-connected CRFs based on low-rank decomposition,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2015.
 119. A. Raj and R. Zabih, “A graph cut algorithm for generalized image deconvolution,” in *Proc. IEEE Int. Conf. Comp. Vis.*, 2005.
 120. J. Malick, “The spherical constraint in boolean quadratic programs,” *J. Global Optimization*, vol. 39, no. 4, pp. 609–622, 2007.