

Submodular Decomposition Framework for Inference in Associative Markov Networks with Global Constraints

Anton Osokin* Dmitry Vetrov*
Department of Computational Mathematics
and Cybernetics, Moscow State University
anton.osokin@gmail.com vetrovd@yandex.ru

Vladimir Kolmogorov
Department of Computer Science
University College London, UK
v.kolmogorov@cs.ucl.ac.uk

Abstract

In this paper we address the problem of finding the most probable state of discrete Markov random field (MRF) with associative pairwise terms. Although of practical importance, this problem is known to be NP-hard in general. We propose a new type of MRF decomposition, submodular decomposition (SMD). Unlike existing decomposition approaches SMD decomposes the initial problem into subproblems corresponding to a specific class label while preserving the graph structure of each subproblem. Such decomposition enables us to take into account several types of global constraints in an efficient manner. We study theoretical properties of the proposed approach and demonstrate its applicability on a number of problems.

The problem of effective Bayesian inference arises in many applied domains, *e.g.* in machine learning, computer vision, decision-making, etc. One of the most intriguing problems is the development of approximate inference algorithms for the problems that are NP-hard in general. An important particular case is the inference problem in cyclic discrete Markov random fields (MRF) with energies that can be represented via sum of unary and pairwise terms.

Let $G = (\mathcal{V}, \mathcal{E})$ be an undirected graph with \mathcal{V} and \mathcal{E} being the sets of nodes and edges respectively. With each node we associate a class label $t_j \in \{1, \dots, P\}$. The inference problem can be formulated as an energy minimization problem

$$\sum_{(j \in \mathcal{V})} \theta_j(t_j) + \sum_{(i,j) \in \mathcal{E}} \theta_{ij}(t_i, t_j) \rightarrow \min_{t_1, \dots, t_{|\mathcal{V}|} \in \{1, \dots, P\}}. \quad (1)$$

where $\theta_j(t_j)$ and $\theta_{ij}(t_i, t_j)$ are some known functions of discrete argument.

Although NP-hard in general there are several special cases when the inference problems can be solved exactly in polynomial time. One example is dynamic programming approach [22] for inference in tree-structured graphs.

Another example are MRFs with outer-planar graphs [24], or more generally MRFs with low treewidth [19]. Min-cut/max-flow algorithms can efficiently solve the inference problem on arbitrary graphs when all the variables are binary and pairwise potentials meet submodularity constraints [8, 15].

$$\theta_{ij}(0, 0) + \theta_{ij}(1, 1) \leq \theta_{ij}(0, 1) + \theta_{ij}(1, 0). \quad (2)$$

The multi-class problems on cyclic graphs do not have efficient exact algorithms except few very special cases with specific orderings in the space of classes [9, 6].

Advanced approximate methods based on MRF decomposition have recently appeared [25]. The most popular method, tree-reweighted message passing (TRW) [29, 13, 17], splits the MRF with cycles into a number of acyclic subgraphs (trees) and, for each tree, inference is made independently with the subsequent harmonization of optimal solutions. The other decomposition methods (*e.g.* [30, 26, 16, 1]) exploit similar ideas. Unlike more efficient approximate energy minimization algorithms, *e.g.* [5, 18], decomposition framework makes it possible to take into account some global properties of the solution, *e.g.* to establish constraints on the areas of classes [31, 32, 20].

In this paper we develop an alternative framework for approximate inference in MRF with associative pairwise terms. We present submodular decomposition (SMD) technique. Instead of subgraph-based decomposition we split the problem into a number of submodular problems keeping the structure of graph \mathcal{G} . To harmonize optimal solutions of different subproblems we use the dual decomposition approach [2]. We show that SMD provides energy lower bounds which are equivalent to the ones by TRW. However the new type of decomposition allows us to take into account the preferences on any type of global linear statistics of the class indicator variables in a straightforward manner as well as to establish some shape constraints.

The rest of paper is organized as follows. In the next section we present SMD framework. We prove the equivalence of SMD and TRW lower bounds in section 2 and examine the properties of SMD convergent point in section 3.

*The authors assert equal contribution and thus joint first authorship.

Section 4 presents a possible way for optimizing the lower bound which is based on the averaging of min-marginals. The different types of global constraints which can be implemented efficiently in our SMD framework are discussed in section 5. We present the experimental evaluation of our method in section 6 and give some conclusions in section 7.

1. Submodular Decomposition

Consider the indicator parametrization of (1) obtained by establishing auxiliary binary variables $Y = \{y_{jp}\} \in \{0, 1\}^{|\mathcal{V}| \times P}$:

$$y_{jp} = \begin{cases} 1, & t_j = p, \\ 0, & \text{otherwise.} \end{cases}$$

Denote $\theta_j(p) = \theta_{jp}$ and $\theta_{ij}(p, q) = \theta_{ij,pq} = \theta_{ji,qp}$. Then the problem (1) takes the form

$$E(Y) = \sum_{j \in \mathcal{V}} \sum_{p=1}^P \theta_{jp} y_{jp} + \sum_{(i,j) \in \mathcal{E}} \sum_{p,q=1}^P \theta_{ij,pq} y_{ip} y_{jq} \rightarrow \min_{Y \in \mathcal{L} \cap \mathcal{G}}. \quad (3)$$

Here we denoted the set of *local* constraints by $\mathcal{L} = \{Y \mid y_{jp} \in \{0, 1\}\}$ and the set of *global* constraints by $\mathcal{G} = \{Y \mid \sum_{p=1}^P y_{jp} = 1, \forall j \in \mathcal{V}\}$. The first group of constraints requires y_{jp} to be indicator variables while the second group provides consistency of graph labeling, i.e. each node belongs to one and only one class.

In what follows we assume the associativeness [27] of pairwise terms, i.e.

$$\theta_{ij,pq} = -C_{ij,p} \delta_{pq}, \quad (4)$$

where $\delta_{pq} = 1$ iff $p = q$ and $C_{ij,p}$ are non-negative constants! Then we may rewrite the energy (3) as follows

$$\begin{aligned} \min_{Y \in \mathcal{L} \cap \mathcal{G}} E(Y) &= \min_{Y \in \mathcal{L} \cap \mathcal{G}} \sum_{p=1}^P \left(\sum_{j \in \mathcal{V}} \left(\theta_{jp} - \frac{1}{2} \sum_{i:(i,j) \in \mathcal{E}} C_{ij,p} \right) y_{jp} + \frac{1}{2} \sum_{(i,j) \in \mathcal{E}} C_{ij,p} \left[y_{ip}(1 - y_{jp}) + y_{jp}(1 - y_{ip}) \right] \right) \\ &= \min_{Y_p \in \{0,1\}^{|\mathcal{V}|}, \sum Y_p = \bar{1}} \sum_{p=1}^P E_p(Y_p) \geq \min_{Y_p \in \{0,1\}^{|\mathcal{V}|}} \sum_{p=1}^P E_p(Y_p) = \sum_{p=1}^P \min_{Y_p \in \{0,1\}^{|\mathcal{V}|}} E_p(Y_p). \quad (5) \end{aligned}$$

Here we denoted the p^{th} column of Y by Y_p and (i, j) is an undirected edge of the graph. Note that each $E_p(Y_p)$ is a submodular function so that we can find $\min_{Y_p \in \{0,1\}^{|\mathcal{V}|}} E_p(Y_p)$ efficiently by a min-cut/max-flow

¹Note that by setting all $C_{ij,p} = C_{ij}$ we get an important case of generalized Potts potential.

algorithm (e.g. [4]). To reconcile the minimizers of E_p consider a Lagrangian

$$L(Y, \Lambda) = \sum_{p=1}^P E_p(Y_p) + \sum_{j \in \mathcal{V}} \lambda_j \left(\sum_{p=1}^P y_{jp} - 1 \right). \quad (6)$$

The following lower bound is valid for the initial problem (3):

$$\begin{aligned} \min_{Y \in \mathcal{L} \cap \mathcal{G}} E(Y) &= \min_{Y \in \mathcal{L} \cap \mathcal{G}} L(Y, \Lambda) \geq \max_{\Lambda} \min_{Y \in \mathcal{L}} L(Y, \Lambda) = \\ &= \max_{\Lambda} \left[\sum_{p=1}^P \min_{Y_p \in \{0,1\}^{|\mathcal{V}|}} \left(E_p(Y_p) + \sum_{j \in \mathcal{V}} \lambda_j y_{jp} \right) - \sum_{j \in \mathcal{V}} \lambda_j \right]. \quad (7) \end{aligned}$$

Note that the minimization of

$$\Phi_p(Y_p, \Lambda) = E_p(Y_p) + \sum_{j \in \mathcal{V}} \lambda_j y_{jp}$$

w.r.t. Y_p can still be done by min-cut algorithms since the functions remain submodular. The maximization problem (7) is concave w.r.t. Λ and hence the solution can be found via subgradient ascent. Thereby we replace the NP-hard optimization problem (3) with a series of simpler problems that can be solved efficiently. Hereinafter we refer to both decomposition and algorithm as submodular decomposition (SMD).

Note, that $\min_{Y \in \mathcal{L} \cap \mathcal{G}} E(Y) = \min_{Y \in \mathcal{L}} \max_{\Lambda} L(Y, \Lambda)$ and energy lower bound (6) equals the maximin value of Lagrangian $\max_{\Lambda} \min_{Y \in \mathcal{L}} L(Y, \Lambda)$. Now consider the continuous relaxation $\mathcal{Q} = \{Y \mid y_{jp} \in [0, 1]\}$ of local constraints. It can be shown [23, Theorem 3.1] that $\min_{Y \in \mathcal{L} \cap \mathcal{G}} E(Y) = \min_{Y \in \mathcal{Q} \cap \mathcal{G}} E(Y)$ for arbitrary number of classes, unary and pairwise potentials, hence

$$\min_{Y_p \in \{0,1\}^{\mathcal{V}}} \Phi_p(Y_p, \Lambda) = \min_{Y_p \in [0,1]^{\mathcal{V}}} \Phi_p(Y_p, \Lambda).$$

This implies

$$\begin{aligned} L^* &= \min_{Y \in \mathcal{L} \cap \mathcal{G}} E(Y) = \min_{Y \in \mathcal{Q}} \max_{\Lambda} L(Y, \Lambda) \geq \\ &= \max_{\Lambda} \min_{Y \in \mathcal{L}} L(Y, \Lambda) = \max_{\Lambda} \min_{Y \in \mathcal{Q}} L(Y, \Lambda) = L^*. \end{aligned}$$

In the space $(\mathcal{Q} \cap \mathcal{G}) \times \mathbb{R}^{|\mathcal{V}|}$ the Lagrangian (6) is a smooth and differentiable function, hence we can find the global minimum of energy (3) by performing SMD if and only if $L(Y, \Lambda)$ has a saddle-point in $(\mathcal{Q} \cap \mathcal{G}) \times \mathbb{R}^{|\mathcal{V}|}$.

2. Equivalence of Lower Bounds

In this section we show that the optimal value of the Lagrangian dual (7), which is solved by SMD, is equivalent Kleinberg-Tardos (KT) relaxation [10] of uniform metric labeling problem. Furthermore we prove that KT-relaxation is equivalent to Schlesinger LP-relaxation which is reviewed in [30] and is solved e.g. by TRW [29].

Theorem 1 The maximin value the Lagrangian (6) equals the minimum value of KT-relaxed problem [10]:

$$L_* = \min_X \sum_{j \in \mathcal{V}} \sum_{p=1}^P \theta_{jp} x_{jp} + \frac{1}{2} \sum_{(i,j) \in \mathcal{E}} \sum_{p=1}^P C_{ij,p} x_{ij,p}, \quad (8)$$

$$\text{s.t. } \sum_{p=1}^P x_{ip} = 1,$$

$$x_{ij,p} \geq x_{ip} - x_{jp}, \quad x_{ij,p} \geq x_{jp} - x_{ip},$$

$$x_{ij,p} \geq 0, \quad x_{ip} \geq 0.$$

See the proof in [21].

Theorem 2 If energy (3) satisfies the condition (4) then the KT-relaxation [10] is equivalent to the Schlesinger LP-relaxation:

$$L_* = \min_Z \sum_{j \in \mathcal{V}} \sum_{p=1}^P \theta_{jp} z_{jp} + \sum_{(i,j) \in \mathcal{E}} \sum_{p,q=1}^P d_{ij,pq} z_{ij,pq}, \quad (9)$$

$$\text{s.t. } \sum_{p=1}^P z_{ip} = 1,$$

$$\sum_{p=1}^P z_{ij,pq} = z_{jq}, \quad \sum_{q=1}^P z_{ij,pq} = z_{ip},$$

$$z_{ij,pq} \geq 0, \quad z_{ip} \geq 0,$$

where $d_{ij,pq} = 0$ iff $p = q$ and $\frac{1}{2}(C_{ij,p} + C_{ij,q})$ otherwise.

See the proof in [21].

3. Properties of the Maximin Point

Consider the properties of $L(Y, \Lambda)$ when $Y \in \mathcal{L}$, i.e. all $y_{jp} \in \{0, 1\}$. Then the maximin problem (7) can be viewed as maximization of a piecewise linear function $\min_{Y \in \mathcal{L}} L(Y, \Lambda)$ in the space of Λ 's. Each $Y \in \mathcal{L}$ defines a hyperplane in this space and $\Omega = \{(\Lambda, l) \mid l \leq \min_{Y \in \mathcal{L}} L(Y, \Lambda)\}$ is a concave polytope.

Lemma 1 Let (Y^0, Λ^0) be a maximin point of the Lagrangian, i.e. $L(Y^0, \Lambda^0) = \max_{\Lambda} \min_{Y \in \mathcal{L}} L(Y, \Lambda)$. The global minimum of energy $E(Y)$ equals $L(Y^0, \Lambda^0)$ if and only if there is a horizontal hyperplane defined by Y^1 such that $L(Y^1, \Lambda^0) = L(Y^0, \Lambda^0)$.

If it is known additionally that the subgradient of $\min_{Y \in \mathcal{L}} L(Y, \Lambda)$ at the point Λ^0 consists of only one vector then Y^0 is the optimal solution of (3).

See the proof in [21].

Since $\min_{Y \in \mathcal{L}} L(Y, \Lambda)$ is a piecewise linear function three cases are possible at its maximum (see fig. 1). In the first case there is a plateau on the top of the polytope. In the second case there is no plateau but there exists a horizontal hyperplane defined by Y^1 that touches polytope Ω . The third case corresponds to the non-zero duality gap. Note that in the first case we get the optimal value of $Y^0 = \arg \min_{Y \in \mathcal{L} \cap \mathcal{G}} E(Y)$ whereas in the second case in

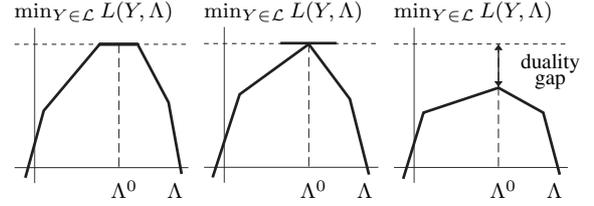


Figure 1. Three possible cases in SMD dual problem. The optimal energy value is shown by dotted line. The left plot corresponds to the case when we are able to find both optimal energy and optimal configuration of labels. The middle plot shows the case when there is zero duality gap but we are not guaranteed to find the optimal configuration of labels. The horizontal face corresponds to Y^1 (see lemma 1). The right plot illustrates the case when there's no horizontal face coming through the maximin point. Then we observe a non-zero duality gap.

general we may not get the optimal value of $Y^1 \in \mathcal{G}$ in an explicit manner. Several cases when it becomes possible are considered further.

Now we define and investigate strong agreement (SA) and weak agreement (WA) conditions that are analogous to strong tree agreement (STA) and weak tree agreement (WTA)².

Denote the possible optimal labelings of the node j in subproblem p given Λ as

$$Z_{jp}(\Lambda) = \{z \mid \exists Y_p^* = (y_{1p}^*, \dots, y_{|\mathcal{V}|p}^*) : y_{jp}^* = z, \Phi_p(Y_p^*, \Lambda) = \min \Phi_p(Y_p, \Lambda)\}.$$

Definition 1 The variables Λ satisfy the strong agreement condition if

$$\forall j \quad \exists! p : Z_{jp}(\Lambda) = \{1\}, \quad \forall q \neq p : Z_{jq}(\Lambda) = \{0\}.$$

From the last definition it follows that Λ^0 satisfies the strong agreement condition if and only if $\arg \min_{Y \in \mathcal{L}} L(Y, \Lambda^0) \in \mathcal{G}$.

Theorem 3 If there exists Λ^0 that satisfies the strong agreement condition and $Y^0 = \arg \min_{Y \in \mathcal{L}} L(Y, \Lambda^0)$, then $E(Y^0) = \min_{Y \in \mathcal{L} \cap \mathcal{G}} E(Y)$.

Definition 2 The variables Λ satisfy weak agreement condition if for all j both statements hold true:

1. $\exists q : 1 \in Z_{jq}(\Lambda)$.
2. If $\exists p : Z_{jp}(\Lambda) = \{1\}$, then $\forall q \neq p \Rightarrow 0 \in Z_{jq}(\Lambda)$;

It is easy to see that strong agreement condition implies weak agreement condition.

Theorem 4 If the point (Λ^0, Y^0) is a maximin point of $L(Y, \Lambda)$ then variables Λ^0 satisfy weak agreement condition.

²WTA and STA were defined in [14] for TRW framework. In this paper we explore similar conditions for SMD.

Proof. Assume the contrary. Let (Λ^0, Y^0) be a maximin point. There are two possible cases: there exist j, p , and q such that $Z_{jp}(\Lambda^0) = Z_{jq}(\Lambda^0) = \{1\}$; there exists j , such that for all p it holds $Z_{jp}(\Lambda^0) = \{0\}$. In the first case we may always take Λ^1 such that for all $i \neq j$ holds $\lambda_i^1 = \lambda_i^0$ and $\lambda_j^1 = \lambda_j^0 + \varepsilon$, where $\varepsilon > 0$. If ε is small enough $\{1\}$ still belongs to both $Z_{jp}(\Lambda^1)$ and $Z_{jq}(\Lambda^1)$. Hence $Y^0 \in \text{Arg min } L(Y, \Lambda^1)$ but $L(Y^0, \Lambda^1) > L(Y^0, \Lambda^0)$ that contradicts the condition that (Λ^0, Y^0) is maximin point of the Lagrangian. The second case is considered in a similar manner. ■

Let $Free(p, \Lambda)$ be a set of nodes such that $Z_{jp}(\Lambda) = \{0, 1\}$, and let $Free^k(p, \Lambda)$ be its k^{th} connected component.

Theorem 5 Let $Y_p^* = \arg \min_{Y_p} \Phi_p(Y_p, \Lambda)$. Consider $Y_p^{(1)}$ such that

$$y_{jp}^{(1)} = \begin{cases} 1, & j \in Free(p, \Lambda), \\ y_{jp}^*, & \text{otherwise}; \end{cases}$$

and $Y_p^{(0)}$ such that

$$y_{jp}^{(1)} = \begin{cases} 0, & j \in Free(p, \Lambda), \\ y_{jp}^*, & \text{otherwise}. \end{cases}$$

Then the following is true

$$\Phi_p(Y_p^{(1)}, \Lambda) = \Phi_p(Y_p^{(0)}, \Lambda) = \Phi_p(Y_p^*, \Lambda).$$

The proof follows from submodularity inequality for the function $E_p(Y_p)$.

Note that the last theorem can be trivially generalized to separate connected components. In each $Free^k(p, \Lambda)$ we may replace y_{jp}^* with either all ones or all zeros regardless of other $Free^l(p, \Lambda)$ and still keep $\Phi_p(Y_p, \Lambda)$ at its minimum. This is an important property since it helps to harmonize the solutions of subproblems.

Theorem 6 Assume that Λ^0 satisfies weak agreement condition. Then for any $\varepsilon > 0$ there always exists Λ^1 such that:

- $\|\Lambda^0 - \Lambda^1\| < \varepsilon$;
- Λ^1 also satisfies weak agreement condition;
- $\min_{Y \in \mathcal{L}} L(Y, \Lambda^0) = \min_{Y \in \mathcal{L}} L(Y, \Lambda^1)$;
- If for any j there is a subproblem p such that $j \in Free(p, \Lambda^1)$ then
 - there exists at least yet another subproblem q such that $j \in Free(q, \Lambda^1)$;
 - for all problems r such that $j \notin Free(r, \Lambda^1)$ it follows $Z_{jr}(\Lambda^1) = \{0\}$.

Proof. We provide here a sketch of the proof.

At the point Λ^0 the Lagrangian has either a plateau or a peak w.r.t. each variable λ_j . In case of plateau if λ_j is on

the border then take $\lambda_j \pm \varepsilon$ in order to be *inside* the plateau. At each interior point of the plateau $Z_{jp}(\Lambda^0)$ consists of a single element for each p .

If there is no plateau then $\frac{\partial L(Y, \Lambda)}{\partial \lambda_j}$ has a jump of at least 2 at the point Λ^0 . This means that there are at least two subproblems p and q such that $Z_{jp}(\Lambda^0) = Z_{jq}(\Lambda^0) = \{0, 1\}$.

If for Λ^0 there exists r such that $Z_{jr} = \{1\}$ then set $\lambda_j^1 = \lambda_j^0 + \varepsilon$. $Z_{jr}(\Lambda^1)$ still contains $\{1\}$ while for all other subproblems $Z_{jq}(\Lambda^1)$ consists of only zeros. ■

Without loss of generality we may consider only those Λ^1 that satisfy the last theorem. Then the following theorems hold.

Theorem 7 Assume that Λ^0 satisfies weak agreement condition and $Y^0 = \arg \min_{Y \in \mathcal{L}} L(Y, \Lambda^0)$. If there exists such subproblem p that for all q it follows that $Free(q, \Lambda^0) \subseteq Free(p, \Lambda^0)$ then setting

$$y_{jr}^* = \begin{cases} 1, & j \in Free(r, \Lambda^0), r = p; \\ 0, & j \in Free(r, \Lambda^0), r \neq p; \\ y_{jr}^0, & \text{otherwise}, \end{cases}$$

we get an optimal configuration $Y^* = \arg \min_{Y \in \mathcal{L} \cap \mathcal{G}} E(Y)$.

This theorem is a corollary of theorems 5 and 6. Note that it can be straightforwardly extended to the case of separate connected components.

Theorem 8 Assume that Λ^0 satisfies weak agreement condition and $Y^0 = \arg \min_{Y \in \mathcal{L}} L(Y, \Lambda^0)$. If there exists a partition $Free^{k_1}(p_1, \Lambda^0), \dots, Free^{k_m}(p_m, \Lambda^0)$ of the set $I = \{j | \exists p : Z_{jp} = \{0, 1\}\}$ such that $Free^{k_i}(p_i, \Lambda^0) \cap Free^{k_j}(p_j, \Lambda^0) = \emptyset$ and $I = \bigcup_i Free^{k_i}(p_i, \Lambda^0)$ then an optimal configuration $Y^* = \arg \min_{Y \in \mathcal{L} \cap \mathcal{G}} E(Y)$ can be constructed as follows

$$y_{jr}^* = \begin{cases} y_{jr}^0, & j \notin I; \\ 1, & r = p_i, j \in Free^{k_i}(p_i, \Lambda^0), \forall i = 1, \dots, m; \\ 0, & \text{otherwise}. \end{cases}$$

Proof. For each i we apply theorem 7 to $Free^{k_i}(p_i, \Lambda^0)$. ■

The last theorem allows us to get an optimal solution of initial problem in the some cases when duality gap is zero but there is no plateau in the top of a polytope defined by $\min_{Y \in \mathcal{L}} L(Y, \Lambda)$ (case 2 in lemma 1).

4. Min-Marginals Averaging

In this section we derive an alternative approach to the optimization of SMD lower bound based on the min-marginals averaging. First we prove an important lemma.

Suppose that for each of our subproblems we have estimated unary min-marginals $\theta_{jp}(k, \Lambda) = \min_{\{Y_p | y_{jp} = k\}} \Phi_p(Y_p, \Lambda)$. This can be done effectively e.g.

using dynamic graphcut framework [11]. Now consider the following re-estimation equations for Lagrange multipliers

$$\lambda_j^{new} = \lambda_j^{old} + \Delta\lambda_j, \quad (10)$$

where $\Delta\lambda_j = \max_p [\hat{\theta}_{jp}(0, \Lambda^{old}) - \hat{\theta}_{jp}(1, \Lambda^{old})]$.

Theorem 9 *The iterative process defined by (10) has the following properties:*

- *The lower bound of the Lagrangian does not decrease at each step;*
- *The fixed point of the process satisfies weak agreement condition.*

See the proof is in [21].

Theorem 10 *If $\Delta\lambda_j^0 = 0$ for some point Λ^0 then Λ^0 is a maximum point of $L(Y, \Lambda)$ w.r.t. λ_j given all other λ 's are fixed.*

Proof. The condition $\Delta\lambda_j = 0$ implies $\max_p(\hat{\theta}_{jp}(0) - \hat{\theta}_{jp}(1)) = 0$. If we increase λ_j all y_{jp} are assigned the value of 0 and hence the Lagrangian decreases. Similarly the decrease of λ_j implies the decrease of the Lagrangian since there is at least one subproblem p for which $y_{jp} = 1$.

From the last theorem it follows that unless Λ^0 satisfies weak agreement condition, among the nodes with $\Delta\lambda_j \neq 0$ there always exists at least one node i such that changing λ_i by $\Delta\lambda_i$ will definitely increase the value of the Lagrangian. The question of how to find such i is easy to answer. Any $\Delta\lambda_j < 0$ leads to the increase of $L(Y, \Lambda)$. If all $\Delta\lambda_j \geq 0$ then we should select i for which there exist at least two subproblems p and q such that

$$\begin{cases} \hat{\theta}_{ip}(0, \Lambda) > \hat{\theta}_{ip}(1, \Lambda) \\ \hat{\theta}_{iq}(0, \Lambda) > \hat{\theta}_{iq}(1, \Lambda). \end{cases}$$

The update of such λ_j increases the value of the Lagrangian³. However we found the process of min-marginals averaging to be too slow in practice even with the use of dynamic graphcuts [12]. So in our implementation we mostly used subgradient ascend switching to min-marginals averaging for one iteration when the lower bound started oscillating.

The possibility of min-marginals averaging and weak agreement condition shows that there are lots of similarities between various types of decomposition and that optimization on trees and optimization of submodular functions on cycled graphs have much in common.

³If it is not possible to make an update according to these rules than we obtain a strong agreement situation.

5. The Inclusion of Global Terms

The submodular decomposition allows us to make use of its form by including several types of the global constraints related to classes. In particular we consider all linear constraints of indicator variables of the following form

$$\sum_{j \in \mathcal{V}} \sum_{p=1}^P w_{jp}^m y_{jp} = c^m, \quad m = 1, \dots, M \quad (11)$$

$$\sum_{j \in \mathcal{V}} \sum_{p=1}^P v_{jp}^k y_{jp} \leq d^k, \quad k = 1, \dots, K. \quad (12)$$

Then p^{th} subproblem takes the form

$$\begin{aligned} \Phi_p(Y_p, \Lambda, M, K) = & E_p(Y_p) + \sum_{j \in \mathcal{V}} \lambda_j y_{jp} + \\ & \sum_{m=1}^M \mu_m \sum_{j \in \mathcal{V}} w_{jp}^m y_{jp} + \sum_{k=1}^K \kappa_k \sum_{j \in \mathcal{V}} v_{jp}^k y_{jp} \rightarrow \min_{Y_p \in \mathcal{L}} \end{aligned}$$

The subproblem is still submodular. The maximization of

$$\begin{aligned} \min_{Y \in \mathcal{L}} L(Y, \Lambda, M, K) = & \sum_{p=1}^P \min_{Y_p} \Phi_p(Y_p, \Lambda, M, K) - \\ & \sum_{j \in \mathcal{V}} \lambda_j - \sum_{m=1}^M \mu_m c^m - \sum_{k=1}^K \kappa_k d^k \end{aligned}$$

w.r.t. Lagrange multipliers Λ , M , and $K \geq \vec{0}$ can be done *e.g.* via subgradient ascent since the function is piecewise linear and concave. Note that each global constraint adds just one additional variable in the Lagrangian.

Here are several examples of global linear constraints:

$$\sum_{j \in \mathcal{V}} y_{jp} = c, \quad \text{Strict class size constraints} \quad (13)$$

$$\sum_{j \in \mathcal{V}} y_{jp} \in [c_1, c_2], \quad \text{Soft class size constraints} \quad (14)$$

$$\sum_{j \in \mathcal{V}} y_{jp} I_j = \sum_{j \in \mathcal{V}} y_{jq} I_j, \quad \text{Flux equalities} \quad (15)$$

$$\begin{cases} \sum_{j \in \mathcal{V}} y_{jp} I_j = \vec{\mu} \sum_{i \in \mathcal{V}} y_{ip}, & \text{Color mean/} \\ \sum_{j \in \mathcal{V}} y_{jp} (I_j - \vec{\mu})(I_j - \vec{\mu})^T = \sum_{j \in \mathcal{V}} y_{ip} \Sigma, & \text{variance.} \end{cases} \quad (16)$$

Here I_j is an observable scalar or vector value associated with node, *e.g.* it can be the color of an image or the intensity of a grayscale image.

Such global linear statistics cannot be used in the state-of-the-art α -expansion [5] method since it does not provide a lower bound for energy and their inclusion in TRW-like algorithms results in the following problem

$$L_{TRW}(Y, \vec{\Theta}) - \sum_m \mu_m c^m - \sum_k \kappa_k d^k \rightarrow \max_{\vec{\Theta}, M, K} \quad (17)$$

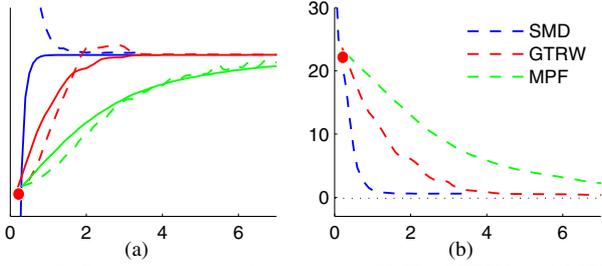


Figure 2. Results on a synthetic test of SMD, GTRW, and MPF. (a): solid lines correspond to lower bounds of different methods, dashed lines show energy of primal solution (that violates global hard constraints) at each iteration; (b): the total constraint violation of primal solution (measured in percentage of pixels that has to be recolored to make constraints consistent). On both plots horizontal axes show time measured in seconds. The red circles on both plots show the result of TRW-S algorithm as a “starting point” of GTRW and MPF.

$$\begin{aligned}
 & s.t. \quad \sum_T \rho^T \tilde{\vartheta}^T \equiv \vartheta \\
 \vartheta_{jp} &= \theta_{jp} + \sum_m \mu_m w_{jp}^m + \sum_k \kappa_k v_{jp}^k \\
 \vartheta_{ij,pq} &= \theta_{ij,pq}, \quad \kappa_k \geq 0.
 \end{aligned}$$

Here T indexes the set of subgraphs that are used in TRW decomposition, $L(Y, \Theta)$ is sum of optimal energies computed on each of the subgraphs, and \equiv defines reparameterization relation. The straightforward way to solve the problem leads to nested iterative process. On the inner loop we maximize L_{TRW} given ϑ and on the outer loop we update M and K by making subgradient ascent. We will refer to this algorithm as GTRW. One could use a linear programming relaxation of (17) but then it’s not clear how to solve it efficiently in the case of general linear constraints such as (13)–(16) (invoking a general-purpose LP solver may be too slow in practice).

Another promising option of SMD is its ability to take into account some shape constraints. In particular we may establish all shape constraints that can be achieved by binary graphcut algorithm, *e.g.* star-shape prior [28] and “boundary position” prior [7] for any number of classes.

6. Experiments

Synthetic problems. We compare the convergence rates of different methods on a number of artificially generated problems with linear constraints. Similarly to [13] as a synthetic setting we take a 10-label problem on a 50×50 grid graph with 4-neighborhood system. Unary potentials are generated as independent gaussians: $\theta_{ip} \sim \mathcal{N}(0, 1)$. Parameters $C_{ij,p} = C_{ij}$ of pairwise potentials are generated as an absolute value of $\mathcal{N}(0, 0.5)$. As global constraints we’ve chosen strict class size constraints (13). Class sizes are deliberately set to be significantly different from the

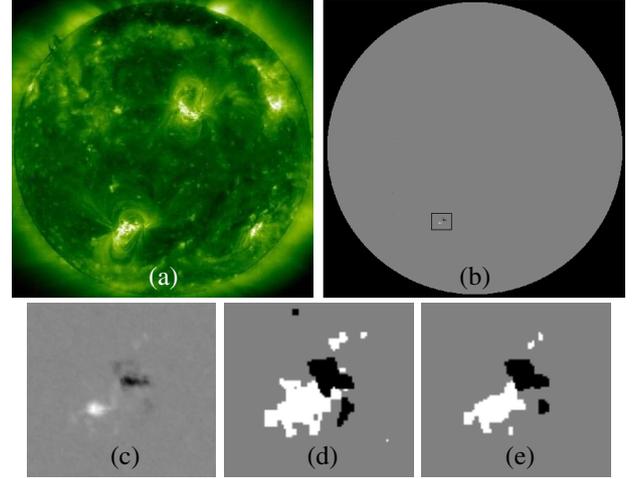


Figure 3. Results on a magnetogram test. (a) – the picture of the Sun at 195 angstrom wavelength in the extreme ultraviolet; (b) – corresponding magnetogram: white regions correspond to the areas with large positive magnetic field, black regions – to large negative magnetic field, gray regions – to relatively inactive areas; (c) – (b) – the increased region of high activity; (d) – the result of α -expansion without global constraints: positive flux = 83794, negative flux = -71021, diff = 12773; (e) – the result of SMD with flux equality constraint: positive flux = 70972, negative flux = -67341, diff = 3630.

global minimum of unconstrained energy (1). Particularly, for class p we set its size to be $|\mathcal{V}|p / \sum_{q=1}^P q$.

Besides SMD we evaluate the performance of GTRW with TRW-S [13] as an inner iterative process. The third method is based on MPF framework proposed in [32]. Let $\Phi_L(\Pi)$ be the lower bound of energy (1) with unary potentials $\vartheta_{jp} = \theta_{jp} - \pi_{jp}$ and pairwise terms $\vartheta_{ij,pq} = \theta_{ij,pq}$; $\Phi_G(\Pi) = \min_Y \langle \Pi, Y \rangle$ s.t. $\sum_j y_{jp} = c_p \geq 0$, $\sum_p c_p = |\mathcal{V}|$, $Y \in \mathcal{Q} \cap \mathcal{G}$. Then we can maximize $\Phi_L(\Pi) + \Phi_G(\Pi)$ w.r.t. Π via subgradient ascent. We computed $\Phi_L(\Pi)$ using TRW-S algorithm and $\Phi_G(\Pi)$ using simplex method for transportation problem.

Figure 2 shows the convergence of SMD, GTRW, and MPF averaged over 50 randomly generated problems. For MPF we do not consider the time required to solve the transportation problem $\Phi_G(\Pi)$ since it is highly dependent on the effectiveness of the implementation. Plot (a) shows lower bounds and energies of primal solutions⁴ of all three methods. (b) shows the reduction of the total constraint violation of the current solution. Note, that energy of primal so-

⁴The choice of primal solution is always performed by an heuristics in dual decomposition methods. In SMD all pixels with conflicting labels in the same connected component were assigned to the same randomly selected label from the set of conflicting labels. In GTRW primal solution was chosen as TRW-S primal solution. In MPF primal solution was chosen as a solution of local subproblem $\Phi_L(\Pi)$ solved by TRW-S (solution of global subproblem $\Phi_G(\Pi)$ always satisfies global constraints but typically has much higher energy).

lution can be lower than the lower bound since it, generally speaking, does not satisfy constraints. Moreover, in GTRW and MPF the energy of the current solution increases since both methods start from energy minimum found by TRW-S and try to make the solution consistent with constraints. Based on figure 2 we conclude that all three methods converge to the same point but in SMD lower bound converges faster and hence primal solution with both low energy and small constraint violation can be found.

Magnetogram segmentation. We demonstrate that SMD can be used for implying flux constraints. Figure 3a, b show photos of the Sun in ultraviolet and magnetogram specters. In the regions of increased sun activity (Figure 3c) the amplitude of magnetic field is significantly larger (large positive magnetic fields are shown in white while large negative fields are shown in black). The greater is the deviation of pixel intensity from the gray level, the larger is the amplitude of magnetic field. The segmentation of sun’s surface into quiet regions and the regions with large positive and negative magnetic fields and the analysis of their mutual allocation is important for further forecast of sun flares. It is known that for stable sunspots the total positive flux approximately equals the total negative flux in some vicinity of a sunspot. This constraint is a particular case of (15) and hence can be taken into account by SMD. Figures 3d–e shows how the segmentation changed after we have added flux constraints to our energy. Note that the major changes happened in the regions where the amplitude of magnetic field was relatively low and therefore those regions could be treated as quiet ones.

Image segmentation. SMD framework gives an opportunity to take into account any constraints that can be worked out with a single binary graphcut. As an example we present our results of including star-shape prior [28] into SMD framework for image segmentation. We work in a framework that can be viewed as a multilabel version of [3], i.e. user defines a small seed region for each class which we use to collect color statistics and to establish the center for star-shape prior.

As unary potentials we use standard $-\log P(I_i | x_i)$ pixel. Color statistics is collected from a small seed region that is provided by user for each class. As pairwise potentials we use generalized Potts Model $C_{ij} = a_1 + a_2 \exp\left(-\frac{(I_i - I_j)^2}{2\sigma^2}\right)$. We use $a_1 = 2$, $a_2 = 20$; for each edge we set σ to be the absolute average intensity difference in the box corresponding pixels. The box size is set to be 20 by 20.

A star-shape prior is a compact shape priors that given a center c makes an object to be star-shaped around the center, i.e. on any ray starting at c the “object” label 1 cannot appear after the “background” label 0. It can be written down via pairwise potentials in the following way:

$$S_{ij} = \begin{cases} 0, & \text{if } f_i = f_j, \\ \infty, & \text{if } f_i = 1 \text{ and } f_j = 0, \\ \beta, & \text{if } f_i = 0 \text{ and } f_j = 1, \end{cases}$$

where j is between c and i . Parameter β corresponds to the “ballooning force” and is set to 0 in our experiments. For more details on star-shape priors see [28].

Figure 4 presents our results. Note that TRW-S and α -expansion converge to similar almost optimal solution but it has poor quality unless we take into account global constraints. This effect is a result of mixed color statistics in unary terms.

7. Discussion and Future Work

In this paper we presented a novel approach to approximate Bayesian inference in associative MRFs with cycles based on the decomposition of initial NP-hard problem into a number of submodular problems. The advantage of such approach is the fact that each subproblem is related to a specific class. Hence any global statistics of classes can be added directly to the subproblems rather than by organizing iterative process which requires full MAP-inference on each iteration. Some questions still remain open:

- Is it possible to derive an efficient algorithm for min-marginal averaging?
- How does the gap between minimax and maximin depend on the number of global constraints?
- Are there any partial label optimality guarantees for partially coincident subproblems solutions?

Two directions of further research work on SMD should be mentioned. We plan to adapt the SMD framework for a wider class of pairwise terms (e.g. for the cases when pairwise terms define semi-metrics in the space of class labels) and to add more elaborate shape priors as global terms.

Acknowledgements This work was supported by the Russian Foundation for Basic Research (projects 10-01-00131 and 10-01-90419), Russian President Grant MK-3827.2010.9, and Royal Academy of Engineering/EPSRC.

References

- [1] D. Batra, A. C. Gallagher, D. Parikh, and T. Chen. Beyond trees: MRF inference via outer-planar decomposition. In *CVPR*, 2010.
- [2] D. Bertsekas. *Nonlinear Programming*. Athena Scientific, 1999.
- [3] Y. Boykov and G. Funka-Lea. Graph Cuts and Efficient N-D Image Segmentation. *IJCV*, 70(2):109–131, 2006.
- [4] Y. Boykov and V. Kolmogorov. An Experimental Comparison of Min-Cut/Max-Flow Algorithms for Energy Minimization in Vision. *PAMI*, 26(9):1124–1137, 2004.

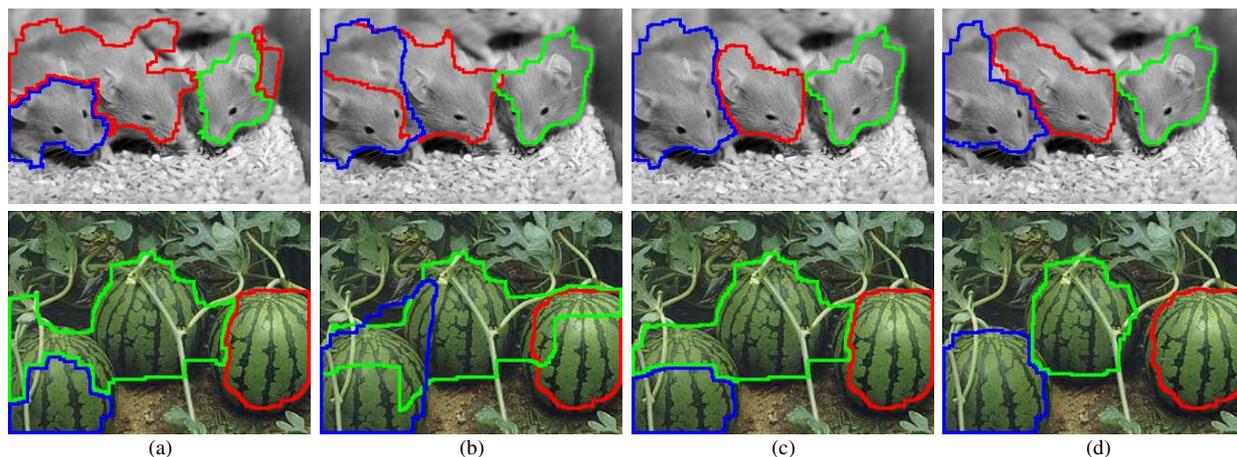


Figure 4. Results for the image segmentation problem. (a) – TRW-S and α -expansion results; (b) – the results of independent usage of binary graphcut with star-shape prior for each object; (c) – the result of SMD with star-shape prior; (d) – the result of SMD with both star-shape prior and class size equality constraints (all classes except background should have equal sizes).

- [5] Y. Boykov, O. Veksler, and R. Zabih. Fast Approximate Energy Minimization via Graph Cuts. *PAMI*, 23(11):1222–1239, 2001.
- [6] J. Darbon. Global optimization for first order Markov random fields with submodular priors. *Discrete Applied Mathematics*, 157(16):3412–3423, 2009.
- [7] D. Freedman and T. Zhang. Interactive Graph Cut Based Segmentation With Shape Priors. In *CVPR*, 2005.
- [8] D. Greig, B. Porteous, and A. Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society, Series B*, 51(2):271–279, 1989.
- [9] H. Ishikawa. Exact Optimization for Markov Random Fields with Convex Priors. *PAMI*, 25(10):1333–1336, 2003.
- [10] J. M. Kleinberg and É. Tardos. Approximation algorithms for classification problems with pairwise relationships: metric labeling and Markov random fields. *J. ACM*, 49(5):616–639, 2002.
- [11] P. Kohli and P. Torr. Measuring Uncertainty in Graph Cut Solutions. In *ECCV*, 2006.
- [12] P. Kohli and P. Torr. Dynamic Graph Cuts for Efficient Inference in Markov Random Fields. *PAMI*, 29(12):2079–1088, 2007.
- [13] V. Kolmogorov. Convergent Tree-Reweighted Message Passing for Energy Minimization. *PAMI*, 28(10):1568–1583, 2006.
- [14] V. Kolmogorov and M. Wainwright. On the Optimality of Tree-reweighted Max-product Message Passing. In *UAI*, 2005.
- [15] V. Kolmogorov and R. Zabih. What Energy Functions Can Be Optimized via Graph Cuts. *PAMI*, 26(2):147–159, 2004.
- [16] N. Komodakis and N. Paragios. Beyond loose LP-relaxations: Optimizing MRFs by repairing cycles. In *ECCV*, 2008.
- [17] N. Komodakis, N. Paragios, and G. Tziritas. MRF Energy Minimization and Beyond via Dual Decomposition. *PAMI*, 33(3):531–552, 2010.
- [18] N. Komodakis and G. Tziritas. Approximate Labeling via Graph-Cuts Based on Linear Programming. *PAMI*, 29(8):1436–1453, 2007.
- [19] S. L. Lauritzen. *Graphical models*. Clarendon Press, 1996.
- [20] Y. Lim, K. Jung, and P. Kohli. Energy Minimization Under Constraints on Label Counts. In *ECCV*, 2010.
- [21] A. Osokin, D. Vetrov, and V. Kolmogorov. Submodular Decomposition Framework for Inference in Associative Markov Networks with Global Constraints. Technical Report 1103.1077, arXiv, 2011.
- [22] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman, San Francisco, 1988.
- [23] P. Ravikumar and J. Lafferty. Quadratic Programming Relaxations for Metric Labeling and Markov Random Fields MAP Estimation. In *ICML*, 2006.
- [24] N. N. Schraudolph and D. Kamenetsky. Efficient Exact Inference in Planar Ising Models. In *NIPS*, 2008.
- [25] D. Sontag, A. Globerson, and T. Jaakkola. Introduction to dual decomposition for inference. In S. Sra, S. Nowozin, and S. J. Wright, editors, *Optimization for Machine Learning*. MIT Press, 2011.
- [26] D. Sontag, T. Meltzer, A. Globerson, Y. Weiss, and T. Jaakkola. Tightening LP relaxations for MAP using message-passing. In *UAI*, 2008.
- [27] B. Taskar, V. Chatalbashev, and D. Koller. Learning Associative Markov Networks. In *ICML*, 2004.
- [28] O. Veksler. Star Shape Prior for Graph-Cut Image Segmentation. In *ECCV*, 2008.
- [29] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky. MAP estimation via agreement on trees: message-passing and linear programming. *IEEE Transactions on Information Theory*, 51(11):3697–3717, 2005.
- [30] T. Werner. A Linear Programming Approach to Max-sum Problem: A Review. *PAMI*, 29(7):1165–1179, 2007.
- [31] T. Werner. High-arity Interactions, Polyhedral Relaxations, and Cutting Plane Algorithm for Soft Constraint Optimisation (MAP-MRF). In *CVPR*, 2008.
- [32] O. J. Woodford, C. Rother, and V. Kolmogorov. A Global Perspective on MAP Inference for Low-Level Vision. In *ICCV*, 2009.