

A discriminative view of MRF pre-processing algorithms

Chen Wang^{1,2} Charles Herrmann² Ramin Zabih^{1,2}
¹ Google Research ² Cornell University
 {chenwang, cih, rdz}@cs.cornell.edu

Abstract

While Markov Random Fields (MRFs) are widely used in computer vision, they present a quite challenging inference problem. MRF inference can be accelerated by pre-processing techniques like Dead End Elimination (DEE) [8] or QPBO-based approaches [18, 24, 25] which compute the optimal labeling of a subset of variables. These techniques are guaranteed to never wrongly label a variable but they often leave a large number of variables unlabeled. We address this shortcoming by interpreting pre-processing as a classification problem, which allows us to trade off false positives (i.e., giving a variable an incorrect label) versus false negatives (i.e., failing to label a variable). We describe an efficient discriminative rule that finds optimal solutions for a subset of variables. Our technique provides both per-instance and worst-case guarantees concerning the quality of the solution. Empirical studies were conducted over several benchmark datasets. We obtain a speedup factor of 2 to 12 over expansion moves [4] without preprocessing, and on difficult non-submodular energy functions produce slightly lower energy.

1. Pre-processing for MRF inference

We address the inference problem for pairwise Markov Random Fields (MRFs) defined over n variables $x = (x_1, \dots, x_n)$, where each x_i is labeled from a discrete label set \mathcal{L} . The MRF can be viewed as a graph $G = (V, E)$ with a neighborhood system $\mathcal{N} : V \mapsto 2^V$. To compute the MAP estimate we minimize the energy

$$E(x) = \sum_{i \in V} \theta_i(x_i) + \sum_{(i,j) \in E} \theta_{ij}(x_i, x_j) \quad (1)$$

where θ_i and θ_{ij} are unary terms and pairwise terms. MRFs are widely used in applications such as image segmentation, stereo, etc [14, 40]. Unfortunately the MRF inference problem is NP-hard even when $|\mathcal{L}| = 2$ (i.e. binary labels) [22].

Many algorithms involve some kind of pre-processing phase that seeks to determine the value of a subset of vari-

ables, thus reducing the complexity of the remaining combinatorial search problem. Pre-processing methods are commonly used in conjunction with graph cuts, a technique that achieves strong performance on both binary and multilabel MRF inference [40]. Graph cuts handle binary MRFs by reduction to min-cut, which is then solved via max-flow (see [3, 9] for reviews). The most widely used graph cut methods for multi-label MRFs are move-making techniques, which generate a new proposal at each iteration and reduce the multi-label problem into a series of binary subproblems (should each variable stick with the old label or switch to the proposed label) and then solved by max-flow/min-cut. Popular algorithms in this family include expansion moves [4] and their generalization to fusion moves [27].

The best known pre-processing methods are Dead End Elimination (DEE) [8] and QPBO [3, 21], but there are a number of others [18, 24, 25, 35, 36, 39, 42]. (Similar approaches are used for other NP-hard problems, a prominent example is Davis-Putnam’s pure literal rule for SAT [7].)

The key weakness of such methods is that they are inherently conservative, since they only label variables whose value can be determined in every global minimum. Yet the MRFs that occur in computer vision are so large that in practice we almost never compute the actual global minimum.¹ As a result, a pre-processing step that is carefully designed to never prune the global minimum is followed by a search step that almost never finds the global minimum. Our fundamental observation is that the pre-processing step can be viewed as a classification problem, and that existing pre-processing methods are designed to avoid false positives (i.e., to never label a variable incorrectly), at the cost of many false negatives (i.e., variables that are left unlabeled). By revisiting this tradeoff we can design techniques where the combination of the pre-processing step and the search step leads to better overall performance, especially on the most difficult problems.

As an example, consider a tiny 8-connected binary MRF with 9 variables (pixels), and suppose we wish to determine by pre-processing that the center pixel should be labeled with 0. In order to soundly compute this by DEE or QPBO,

¹See [12, 29] for rare counterexamples.

1	0	1
0	?	0
1	0	1

1	1	1
1	?	1
1	1	1

Table 1. Minimal MRF example. The left configuration is an unlikely configuration for the neighborhood around the center pixel in the global minimum, compared to the configuration at right. Existing pre-processing methods treat all configurations equally, and as a result fail to label many variables.

we need to establish that switching the center pixel from 1 to 0 will always decrease the energy, no matter what the configuration of the surrounding pixels. Yet as demonstrated in Table 1, there are local configurations that are quite unlikely.

1.1. Outline and contributions

We begin with a summary of related work, with an emphasis on DEE, QPBO and QPBO-based pre-processing techniques. In Section 3 we give our discriminative criterion for pre-processing, motivated by examples like Table 1, and provide efficient approximations for the key sub-problems. The theoretical performance of our method is analyzed in Section 4, and experimental results are given in Section 5. Most proofs are deferred to the supplemental material, which also contains additional experimental results.

2. Related work

A popular approach to the inference problem is to find the optimal labeling for a subset of the variables [8, 13, 18, 20, 34, 35, 36, 39, 44]. A partial labeling that holds in every global minimizer is said to be *persistent* [3]. Techniques like QPBO [3, 21] find an optimal partial labeling by enforcing an even stronger condition: a partial labeling that will decrease the energy if it is substituted into any complete labeling.² This stronger property is sometimes called an *autarky* [3], which was generalized by [35]. QPBO in particular is widely used in computer vision since it often finds the correct label for the majority of the variables.

To make these notions precise, we introduce the following notation. A *partial labeling* x_S is represented by the subvector of x indexed by $S \subseteq V$. Let $\mathcal{L}_S = \prod_{i \in S} \mathcal{L}_i$ be the label space of x_S . Given two partial labelings x_A and x_B where $A \cap B = \emptyset$, we define $x_A \oplus x_B$ to be the composition of x_A and x_B .³ As an important special case, we can write substituting a partial labeling x_S into a full labeling z as $x_S \oplus z_{V \setminus S}$.

Following [3], we can define *persistence* and *autarky*:

²QPBO is naturally viewed as a pre-processing method since it finds persistent partial labelings, and leaves the task of labeling the remaining variables to some other algorithm.

³Let $y = x_A \oplus x_B$ when $A \cap B = \emptyset$, then we have $y_i = (x_A)_i$ when $i \in A$ and $y_i = (x_B)_i$ when $i \in B$.

Definition 1. A partial labeling x_S is persistent if

$$x_S = x_S^*, \quad \forall x^* \in \arg \min_x E(x). \quad (2)$$

Definition 2. A partial labeling x_S is an autarky if

$$E(x_S \oplus z_{V \setminus S}) < E(y_S \oplus z_{V \setminus S}), \\ \forall z_{V \setminus S} \in \mathcal{L}_{V \setminus S} \text{ and } \forall y_S \in \mathcal{L}_S \text{ where } y_S \neq x_S. \quad (3)$$

Persistence is the key property for pre-processing, since it determines the optimal value of a subset of the variables and thus reduces the remaining combinatorial search problem. In general, though, checking for persistence is intractable [3]. All existing persistence algorithms appear to check the autarky property as a sufficient condition, which states that overwriting an arbitrary labeling with this partial labeling will reduce the energy.

2.1. MRF pre-processing algorithms

QPBO generalizes the binary graph cut reduction that uses max-flow to find an optimal partial labeling [3, 21, 33]. If the energy function is submodular⁴ the partial labeling is complete (i.e., it labels every variable and finds a global minimizer). However, the computational expense of running max-flow is non-trivial.

There are also techniques directly finding optimal partial labeling for the multi-label case, but the computational costs for these methods are significant. Kovtun [24, 25] described an approach constructing a series of binary one-verse-the-rest auxiliary problems and solve each of them via graph cuts. MQPBO [18] and generalized roof duality [44] proposed generalizations of QPBO to multi-label MRFs.

Recently, Swoboda et. al. [39] use standard MRF inference algorithms to iteratively update the set of persistent variables. Shekhovtsov [35] formalized the problem to maximize the number of optimally labeled variables as an LP. They also proposed to combine these two approaches together which can take advantage of both of them [36]. The number of variables labeled by these approaches are significantly more than Kovtun’s approach and MQPBO. However, the running time of these approaches is significantly longer, since these approaches involve solving complex programming (either via standard MRF inference solver or LP solver) iteratively.

Dead End Elimination (DEE) [8] and the recent Persistence Relaxation (PR) algorithm [42] are the only existing method with cheaper computational costs than max-flow. DEE checks a local sufficient condition which only involves a single vertex and its adjacent edges. PR generalizes DEE to check a larger partial labeling, which gives improved results on standard benchmarks.

⁴For every pairwise cost, we have $\theta_{ij}(0, 0) + \theta_{ij}(1, 1) \leq \theta_{ij}(0, 1) + \theta_{ij}(1, 0)$.

Methods that optimally label a subset of the variables can obviously be used to pre-process and accelerate MRF inference algorithms such as expansion moves. For example, Radhakrishnan and Su [32] used DEE while Alahari et. al. [2] applied Kovtun’s approach.

3. Discriminative pre-processing of MRFs

In computer vision, the MRF inference problem is almost never solved exactly. As a result, pre-processing methods that enforce soundness are far too conservative, since they leave a large number of variables unlabeled. If we view pre-processing as a binary classification problem (given a partial labeling x_S , decide if it’s persistent), existing techniques ensure that there are no false positives (i.e., variables given a label must be part of every global minimum), but at the cost of multiple false negatives (i.e., variables that are left unlabeled).

First, we need some notation. Define

$$\Delta E(y_S \leftarrow x_S \mid z_{V \setminus S}) = E(y_S \oplus z_{V \setminus S}) - E(x_S \oplus z_{V \setminus S})$$

to be the energy change when we substitute x_S by y_S given the partial labeling $z_{V \setminus S}$ for the variables not in S . By expanding the definition of $E(x)$ and cancelling terms, the Markov property of MRFs gives us a sum over terms only depending on x_i, y_i for $i \in S$ and z_j for $j \in V \setminus S$ with some $i \in S$ such that $(i, j) \in E$ (i.e., z_j is adjacent to S). Let $\mathcal{N}(S) = \{j \in V \setminus S \mid \exists i \in S, (i, j) \in E\}$, and we can rewrite $\Delta E(y_S \leftarrow x_S \mid z_{V \setminus S}) = \Delta E(y_S \leftarrow x_S \mid z_{\mathcal{N}(S)})$.

This allows us to rewrite the autarky property (3) as:

$$\min_{y_S \neq x_S} \Delta E(y_S \leftarrow x_S \mid z_{\mathcal{N}(S)}) > 0, \forall z_{\mathcal{N}(S)} \in \mathcal{L}_{\mathcal{N}(S)} \quad (4)$$

The key issue is the universal quantification in Eq. 4. To ensure that a partial labeling x_S presents in all global minimizer, we look at all possible values that the neighbors might have. For each of these, we check that any other assignment y_S would increase the energy.

Yet this is obviously quite conservative. We now show the desired persistency property can be rewritten by only looking at assignments to the neighboring variables that occur in a global minimizer. Define $\mathcal{L}_{\mathcal{N}(S)}^* = \{z_{\mathcal{N}(S)}^* \mid z^* \in \arg \min E(z)\}$ be all possible configurations of $\mathcal{N}(S)$ in a global minimizer.

Lemma 3. x_S is persistent if and only if

$$\min_{y_S \neq x_S} \Delta E(y_S \leftarrow x_S \mid z_{\mathcal{N}(S)}) > 0, \forall z_{\mathcal{N}(S)} \in \mathcal{L}_{\mathcal{N}(S)}^*. \quad (5)$$

Proof. The if direction is trivial: consider an arbitrary global minimizer z^* , we have $z_{\mathcal{N}(S)}^* \in \mathcal{L}_{\mathcal{N}(S)}^*$ by definition. Suppose $x_S \neq z_S^*$, we will have $E(x_S \oplus z_{V \setminus S}^*) < E(z^*)$,

which contradicts the assumption that z^* is a minimizer. Therefore, we have $x_S = z_S^*, \forall z^*$, so it is persistent.

For the only if direction, suppose Eq. 5 is not true, then $\exists z_{\mathcal{N}(S)} \in \mathcal{L}_{\mathcal{N}(S)}^*, \exists y_S \neq x_S$ such that $\Delta E(y_S \leftarrow x_S \mid z_{\mathcal{N}(S)}) \leq 0$. We can expand $z_{\mathcal{N}(S)}$ to one minimizer z^* such that $z_{\mathcal{N}(S)}^* = z_{\mathcal{N}(S)}$. Since x_S is persistent, we also know $z_S^* = x_S$. Therefore, $E(y_S \oplus z_{V \setminus S}^*) \leq E(x_S \oplus z_{V \setminus S}^*) = E(z^*)$. Since z^* is a minimum this inequality is an equality, hence $y_S \oplus z_{V \setminus S}^*$ is also a global minimum. This contradicts the assumption that x_S is persistent, since $y_S \neq x_S$. \square

3.1. Discriminative criterion

Comparing Eq. 4 and Eq. 5, we immediately observe that the universal quantifier makes autarky a sound but stronger condition than persistency. Crucially, this suggests a discriminative criterion to trade off false positives against false negatives.

The high level idea is the following. Let $\hat{\mathcal{L}}_{\mathcal{N}(S)}(x_S) = \{z_{\mathcal{N}(S)} \in \mathcal{L}_{\mathcal{N}(S)} \mid \min_{y_S \neq x_S} \Delta E(y_S \leftarrow x_S \mid z_{\mathcal{N}(S)}) > 0\}$ be the set of neighbor configurations $z_{\mathcal{N}(S)}$ such that given them x_S is always a better choice. When $\hat{\mathcal{L}}_{\mathcal{N}(S)}(x_S)$ is large enough or covers the most important neighbor configurations, it’s very likely that we will have $\mathcal{L}_{\mathcal{N}(S)}^* \subseteq \hat{\mathcal{L}}_{\mathcal{N}(S)}(x_S)$. This in turn implies x_S is persistent, even though $\hat{\mathcal{L}}_{\mathcal{N}(S)}(x_S) \neq \mathcal{L}_{\mathcal{N}(S)}$ and we do not precisely know $\mathcal{L}_{\mathcal{N}(S)}^*$.

Formally, assume we have a ground truth distribution $p(z_{\mathcal{N}(S)})$ which is uniform over $\mathcal{L}_{\mathcal{N}(S)}^*$ and 0 otherwise. Then a sound condition to check persistency is $\sum_{z_{\mathcal{N}(S)} \in \hat{\mathcal{L}}_{\mathcal{N}(S)}(x_S)} p(z_{\mathcal{N}(S)}) = 1$. Of course, computing $\mathcal{L}_{\mathcal{N}(S)}^*$ and $p(z_{\mathcal{N}(S)})$ is computationally intractable. So we use an estimated distribution $q(z_{\mathcal{N}(S)})$ that approximates $p(z_{\mathcal{N}(S)})$. Looking back to Table 1, one would assume that the left configuration would not appear in $Z_{\mathcal{N}(S)}^*$, while the right one quite plausibly could; there should be a lower q value for the left one but a higher q value for the right. Our discriminative criterion for persistency is

$$\sum_{z_{\mathcal{N}(S)} \in \hat{\mathcal{L}}_{\mathcal{N}(S)}(x_S)} q(z_{\mathcal{N}(S)}) \geq \kappa. \quad (6)$$

Here $\kappa \in [0, 1]$ is the key parameter that controls the trade-off between false positives and false negatives, as shown by the following (obvious) lemma.

Lemma 4. For the same set of decision problems for persistency, we will never increase the number of false positives by increasing κ .

We now address the two crucial issues: how to choose q to effectively approximate p , and how to efficiently check Eq. 6.

3.2. Approximating p

A trivial baseline is to treat each $z_{\mathcal{N}(S)}$ as equally important and set our approximation $q(z_{\mathcal{N}(S)})$ to be the uniform distribution over $\mathcal{L}_{\mathcal{N}(S)}$. In this special case, Eq. 6 is equivalent to count the number of neighbor configurations $z_{\mathcal{N}(S)}$ that satisfy $\min_{y_S \neq x_S} \Delta E(y_S \leftarrow x_S | z_{\mathcal{N}(S)}) > 0$. We expect $\hat{\mathcal{L}}_{\mathcal{N}(S)}(x_S)$ to cover the unknown $\mathcal{L}_{\mathcal{N}(S)}^*$ with high probability when $|\hat{\mathcal{L}}_{\mathcal{N}(S)}(x_S)|$ is large enough.

A more elegant approach is to estimate the marginal probability of a particular assignment $z_{\mathcal{N}(S)}$ via the generative MRF model, and use this as our approximation for p . This problem is well studied in the message passing literature, and is often solved by max-product loopy belief propagation (LBP) [31, 43]. An important special case is if we only use the initialization of LBP, $q_i(z_i) \propto e^{-\theta_i(z_i)}$. This makes a certain amount of intuitive sense: in the MRF energy functions that occur in computer vision it is well known that most of the weight comes from the unary terms [31], which provide a strong signal as to the optimal label for each variable.

More generally, we can define $q(z_{\mathcal{N}(S)})$ to be a fully independent distribution $q(z_{\mathcal{N}(S)}) = \prod_{i \in \mathcal{N}(S)} q_i(z_i)$ with $q_i(z_i) \propto e^{-\theta_i(z_i)} \prod_{j \in \mathcal{N}(i)} m_{j \rightarrow i}(z_i)$, where $m_{j \rightarrow i}(z_i)$ is the message we have from the belief propagation algorithm. Since this is just an approximation, we would not need to pay the cost of running LBP to convergence. In our experiments, the more general approach does not seem to pay dividends, but other ways of estimating the marginals are worth investigating.

3.3. Efficiently checking our discriminative criterion

Checking Eq. 6 is generally computational intractable, due to the size of $\mathcal{L}_{\mathcal{N}(S)}(x_S)$ and $\{y_S \in \mathcal{L}_S | y_S \neq x_S\}$. We now propose a polynomial time algorithm to compute a lower bound for $\sum_{z_{\mathcal{N}(S)} \in \hat{\mathcal{L}}_{\mathcal{N}(S)}(x_S)} q(z_{\mathcal{N}(S)})$.

We will focus on the persistency of a single variable x_i from this point forward. This subroutine is used by our construction algorithm (which will be described in Section 3.4) to construct a partial labeling for the given energy function $E(x)$. However, our methods can handle an arbitrary x_S for $|S| > 1$; the details are deferred to the supplementary material, but are similar to the single variable case.

Our general strategy is to find a subset of $\mathcal{L}_{\mathcal{N}(i)}$ which we know is inside $\hat{\mathcal{L}}_{\mathcal{N}(i)}(x_i)$ and can be easily factorized. We start by considering each node $j \in \mathcal{N}(i)$ independently. For each j , define \mathcal{A}_j to be the set of labels ℓ where the autarky condition holds if $z_j = \ell$. Since autarky is a stronger condition than persistency, we know that all $z_{\mathcal{N}(i)}$ values where $z_j \in \mathcal{A}_j$ are inside $\hat{\mathcal{L}}_{\mathcal{N}(i)}(x_i)$. The union of these sets across different $j \in \mathcal{N}(i)$ will still be a subset of $\hat{\mathcal{L}}_{\mathcal{N}(i)}(x_i)$.

Formally, define $\mathcal{L}_{\mathcal{N}(i)}^{z_j = \ell} = \{z_{\mathcal{N}(i)} | z_j = \ell\}$. Then $\mathcal{A}_j = \{\ell | \min_{y_i \neq x_i} \Delta E(y_i \leftarrow x_i | z_{\mathcal{N}(i)}) > 0, \forall z_{\mathcal{N}(i)} \in \mathcal{L}_{\mathcal{N}(i)}^{z_j = \ell}\}$. Let $\mathcal{L}_{\mathcal{N}(i)}^{z_j \in \mathcal{A}_j} = \cup_{\ell \in \mathcal{A}_j} \mathcal{L}_{\mathcal{N}(i)}^{z_j = \ell}$. Then, we know that $\mathcal{L}_{\mathcal{N}(i)}^{z_j \in \mathcal{A}_j} \subseteq \hat{\mathcal{L}}_{\mathcal{N}(i)}(x_i)$ and $\cup_{j \in \mathcal{N}(i)} \mathcal{L}_{\mathcal{N}(i)}^{z_j \in \mathcal{A}_j} \subseteq \hat{\mathcal{L}}_{\mathcal{N}(i)}(x_i)$.

We establish a computationally tractable lower bound for $\sum_{z_{\mathcal{N}(i)} \in \hat{\mathcal{L}}_{\mathcal{N}(i)}(x_i)} q(z_{\mathcal{N}(i)})$ by the following lemma, which we can check instead.

Lemma 5. *We have the following lower bound:*

$$\sum_{j \in \mathcal{N}(i)} Q_j \prod_{k \in \mathcal{N}(i), k \prec j} (1 - Q_k) \leq \sum_{z_{\mathcal{N}(i)} \in \hat{\mathcal{L}}_{\mathcal{N}(i)}(x_i)} q(z_{\mathcal{N}(i)}), \quad (7)$$

where $Q_i = \sum_{\ell \in \mathcal{A}_i} q_i(z_i = \ell)$.

Proof. We can view $\sum_{z_{\mathcal{N}(i)} \in \mathcal{L}'_{\mathcal{N}(i)}} q(z_{\mathcal{N}(i)})$ as the probability $Pr(z_{\mathcal{N}(i)} \in \mathcal{L}'_{\mathcal{N}(i)})$ given distribution q .

Because our $q(z_{\mathcal{N}(i)})$ can be factorized independently, we can integrate over the variables other than z_j to get $Pr(z_{\mathcal{N}(i)} \in \mathcal{L}_{\mathcal{N}(i)}^{z_j = \ell}) = Pr(z_j = \ell) = q_j(z_j = \ell)$.

We also have $Pr(z_{\mathcal{N}(i)} \in \mathcal{L}_{\mathcal{N}(i)}^{z_j \in \mathcal{A}_j}) = Pr(z_j \in \mathcal{A}_j) = \sum_{\ell \in \mathcal{A}_j} q_j(z_j = \ell) = Q_j$ since $\mathcal{L}_{\mathcal{N}(i)}^{z_j = \ell}$ are all disjoint. Then, using independence again, we have

$$\begin{aligned} & \sum_{z_{\mathcal{N}(i)} \in \cup_{j \in \mathcal{N}(i)} \mathcal{L}_{\mathcal{N}(i)}^{z_j \in \mathcal{A}_j}} q(z_{\mathcal{N}(i)}) \\ &= Pr\left(z_{\mathcal{N}(i)} \in \cup_{j \in \mathcal{N}(i)} \mathcal{L}_{\mathcal{N}(i)}^{z_j \in \mathcal{A}_j}\right) \\ &= Pr\left(\cup_{j \in \mathcal{N}(i)} (z_{\mathcal{N}(i)} \in \mathcal{L}_{\mathcal{N}(i)}^{z_j \in \mathcal{A}_j})\right) \\ &= Pr\left(\cup_{j \in \mathcal{N}(i)} (z_j \in \mathcal{A}_j)\right) \\ &= Pr(z_{j_1} \in \mathcal{A}_{j_1}) + Pr(z_{j_2} \in \mathcal{A}_{j_2}) Pr(z_{j_1} \notin \mathcal{A}_{j_1}) \cdots \\ &= \sum_{j \in \mathcal{N}(i)} Q_j \prod_{k \in \mathcal{N}(i), k \prec j} (1 - Q_k) \end{aligned} \quad (8)$$

Finally, note that we argued $\cup_{j \in \mathcal{N}(i)} \mathcal{L}_{\mathcal{N}(i)}^{z_j \in \mathcal{A}_j} \subseteq \hat{\mathcal{L}}_{\mathcal{N}(i)}(x_i)$ before, which concludes the proof. \square

Constructing \mathcal{A}_j requires us to be able to efficiently check $\min_{y_i \neq x_i} \Delta E(y_i \leftarrow x_i | z_{\mathcal{N}(i)}) > 0, \forall z_{\mathcal{N}(i)} \in \mathcal{L}_{\mathcal{N}(i)}^{z_j = \ell}$. We expand it by the definition of $E(x)$ then swap the min and sum operators. This gives the following lower bound, which we check for being strictly positive:

$$\begin{aligned} & \min_{y_i \neq x_i} (\theta_i(y_i) - \theta_i(x_i)) + \min_{y_i \neq x_i} (\theta_{ij}(y_i, \ell) - \theta_{ij}(x_i, \ell)) \\ & + \sum_{k \in \mathcal{N}(i), k \neq j} \min_{z_k, y_i \neq x_i} (\theta_{ij}(y_i, z_k) - \theta_{ij}(x_i, z_k)) > 0 \end{aligned} \quad (9)$$

Algorithm 1: MRF inference with pre-processing	
Input: Energy function $E(x)$	
1	$\hat{x} \leftarrow \emptyset; \quad S \leftarrow \emptyset;$
2	for $t \leftarrow 1$ to τ do
3	for $i \in V \setminus S, \ell \in \mathcal{L}_i$ do
4	Compute $LB \leq \sum_{z_{\mathcal{N}(i)} \in \hat{\mathcal{L}}_{\mathcal{N}(i)}(x_i=\ell)} q(z_{\mathcal{N}(i)});$
5	if $LB \geq \kappa$ then
6	$\hat{x} \leftarrow \hat{x} \oplus \{x_i = \ell\};$
7	$\mathcal{L}_i \leftarrow \{\ell\}; \quad S \leftarrow S \cup \{i\};$
8	end
9	end
10	end
11	With \hat{x}_S fixed, use one MRF inference algorithm to solve the remaining variables, get $\hat{x}_{V \setminus S};$
12	return $\hat{x} = \hat{x}_S \oplus \hat{x}_{V \setminus S};$

3.4. Our algorithm

We have presented our discriminative criterion to decide if a given partial labeling $x_i = \ell$ is persistent. Now we will use it as a key subroutine to perform pre-processing for MRF inference, as shown in lines 2-10 of Algorithm 1. We firstly loop over the unlabeled variables and its label set (line 3). For each given $x_i = \ell$, use our discriminative rule to judge if it's persistent (line 4-8). We will fix its value if it passes our test by setting $\mathcal{L}_i = \{\ell\}$, and concatenate it with our inference result \hat{x} (line 6, 7). Note that fixing $x_i = \ell$ will also provide additional information as to the unlabeled variables which were checked before x_i , so we repeat the whole procedure for τ iterations (line 2).

After our pre-processing has terminated and labeled the variables in the set S , we fix the variables \hat{x}_S and use any MRF inference algorithms to solve the remaining energy minimization problem, which gives us a labeling $\hat{x}_{V \setminus S}$ on the remaining variables (line 11). Finally, we obtain our inference result by concatenating them together (line 12).

Running time We now give an asymptotic bound on the running time of our pre-processing algorithm here, deferring the analysis into the supplementary material. Assuming we have an oracle to give us data terms $\theta_i(x_i)$ and prior terms $\theta_{ij}(x_i, x_j)$ in $\mathcal{O}(1)$ time. Let $N = |V|, M = |E|$ and $L = \max_i |\mathcal{L}_i|$ be the number of variables, edges and maximum possible labels, and $d = \max_i |\mathcal{N}(i)|$ be the maximum degree of the graph. The overall running time is $\mathcal{O}(d^2 NL^2 + EL^2)$ when we use Section 3.3 to check our discriminative criterion, and $\mathcal{O}(dNL^{d+2})$ for brute force (which is feasible when both d and L are small constants).

4. Performance bounds

We can analyze the per-instance and worst-case performance of our pre-processing methods when followed by an

inference algorithm that produces a solution with performance bounds.

4.1. Per-instance bounds

There are a number of MRF inference algorithms that produce per-instance guarantees (i.e., they produce a certificate after execution that their solution is close to the global minimum). These methods, which are typically based on linear programming, include [19, 23, 41], and they provide a per-instance additive error bound by computing the dual-gap.

Our algorithm has a natural way to bound additive errors. Recall our notation $\Delta E(y_i \leftarrow x_i \mid z_{\mathcal{N}(i)})$ describing the energy changes when we flip x_i to y_i with the neighbor configuration $z_{\mathcal{N}(i)}$. Therefore, $\min_{z_{\mathcal{N}(i)}} \min_{y_i} \Delta E(y_i \leftarrow x_i \mid z_{\mathcal{N}(i)}) \leq 0$ is the worst case energy decrement when we flip x_i to arbitrary y_i with arbitrary neighbor configurations $z_{\mathcal{N}(i)}$. It's non-positive since we can always set $y_i = x_i$. Now we can negate it and define $\delta_i \triangleq -\min_{z_{\mathcal{N}(i)}} \min_{y_i} \Delta E(y_i \leftarrow x_i \mid z_{\mathcal{N}(i)})$ to be the maximum potential energy loss when we use our discriminative criterion to decide x_i is persistent. Then we have the following two lemmas.

Lemma 6. *Let \hat{x}_S be the persistent variables found by our Algorithm 1. For arbitrary $\hat{x}_{V \setminus S}$, and arbitrary x'_S , we have $E(\hat{x}_S \oplus \hat{x}_{V \setminus S}) \leq E(x'_S \oplus \hat{x}_{V \setminus S}) + \sum_{i \in S} \delta_i$.*

Proof. With $\hat{x}_{V \setminus S}$ fixed, we flip \hat{x}_i to x'_i in the reverse order of them being added to S by our algorithm. Due to the analysis before, we will lose at most δ_i at each step. \square

Theorem 7. *Suppose the inference algorithm has per-instance ζ -additive bound, then $E(\hat{x}) \leq E(x^*) + \zeta + \sum_{i \in S} \delta_i$.*

Proof. Let $\bar{x}_{V \setminus S}$ be the minimizer of $E(x)$ with \hat{x}_S fixed, which might be different than the global minimizer $x_{V \setminus S}^*$. Then we will have $E(\hat{x}_S \oplus \hat{x}_{V \setminus S}) \leq E(\hat{x}_S \oplus \bar{x}_{V \setminus S}) + \zeta \leq E(\hat{x}_S \oplus x_{V \setminus S}^*) + \zeta \leq E(x_S^* \oplus x_{V \setminus S}^*) + \zeta + \sum_{i \in S} \delta_i$. The first step is because we use an inference algorithm with ζ -additive errors to solve the problem with \hat{x}_S fixed. The second step follows because $\bar{x}_{V \setminus S}$ is the minimizer w.r.t. \hat{x}_S . \square

As a special case, any sound condition like Eq. 4 guarantees $\delta_i = 0$, i.e., we don't make mistakes. In practice it is computationally intractable to compute δ_i , so just as in Section 3.3 we swap the min and sum operators, and compute the upper bound $\bar{\delta}_i \geq \delta_i$ efficiently. Then we use $\sum_{i \in S} \bar{\delta}_i$ as our per-instance additive bound.

4.2. Worst case bounds

Some MRF inference algorithms produce a solution that is guaranteed to lie within a known factor of the global minimum. The best known such technique is the expansion

move algorithm [4] but there are others [11, 17, 23]. We can easily turn our per-instance bounds into the worst case bounds. We combine Eq. 6 and $\bar{\delta}_i \leq \epsilon$ as our discriminative criterion for pre-defined ϵ .

Corollary 8. *Suppose the inference algorithm has worst case ζ -additive bound, then $E(\hat{x}) \leq E(x^*) + \zeta + |S|\epsilon$ is our worst case additive bound.*

Inference algorithm with worst case guarantees are usually multiplicative bounds other than additive bounds, but we can modify our proof of Theorem 7 to get the following bounds.

Theorem 9. *Suppose the inference algorithm has a worst case β -multiplicative bound, then we will have $E(\hat{x}) \leq \beta \cdot E(x^*) + \beta \cdot |S|\epsilon$.*

Proof. Following the proof of Theorem 7, we have $E(\hat{x}_S \oplus \hat{x}_{V \setminus S}) \leq \beta \cdot E(\hat{x}_S \oplus \bar{x}_{V \setminus S}) \leq \beta \cdot E(\hat{x}_S \oplus x_{V \setminus S}^*) \leq \beta \cdot (E(x_S^* \oplus x_{V \setminus S}^*) + |S|\epsilon)$. \square

A more careful analysis can give us a tighter bound (dropping the coefficient β before $|S|\epsilon$), for the important special case where we use the expansion move algorithm [4] for inference. We defer the proof to the supplementary material.

Theorem 10. *Suppose we use expansion moves as the inference algorithm, with the β -multiplicative bound, then we will have $E(\hat{x}) \leq \beta \cdot E(x^*) + |S|\epsilon$.*

5. Experiments

5.1. Datasets and experimental setup

Approaches The most natural baselines for us to compare against include inference without pre-processing, and inference using the sound (but conservative) DEE [8] and PR [42] techniques. We employ expansion moves for MRF inference [4]. In order to achieve better speedup, we apply preprocessing to each induced binary subproblem of expansion moves as the input of DEE, PR or Algorithm 1, and then run QPBO [21].

At the other end of the spectrum are high overhead techniques such as Kovtun’s approach [24, 25], MQPBO [18], and LP-based approaches [35, 36, 39]. These algorithms require more running time than max-flow on each induced binary subproblem. Therefore, we apply them to the multi-label problem, and then use expansion move to infer the remaining part. We choose the IRI method [36] as the representative among [35, 36, 39] since it’s significantly faster. Note that the R³ [2] method also uses Kovtun’s method as their pre-processing (reduce) step in order to speed up MRF inference. The reuse and recycle parts attempt to speed up the inference algorithm itself, which is orthogonal to what

we propose to do in this paper, so we do not compare against this method.

We also compared against other widely used MRF inference algorithms besides expansion moves, including loopy belief propagation (LBP) [31, 43], dual decomposition (DD) [15], TRWS [19] and MPLP [10, 37, 38]. The comparison among these inference algorithms are provided in survey papers [14, 40]. In our experiments, expansion moves is usually significantly faster than other methods, and gives comparable or better energy. These experimental comparisons are deferred to the supplemental material.

Dataset We conducted experiments on a variety of MRF inference benchmarks, where the energy minimization problems come from different vision problems, including color segmentation [26], stereo, image inpainting, denoising [40] and optical flow [6]. Datasets for the first three tasks are wrapped in OpenGM2 [14] and are available online. We use the BSDS300 [28] for the denoising task with the MRF setup following [40]. We use the MPI Sintel dataset [5] for the optical flow task with the MRF setup following [6].

Our focus, of course, is on the difficult inference problems where the induced binary subproblem is non-submodular. For comparison, we also included some experiments on relatively easy problems where the induced binary subproblem is submodular.

Measurement We report the improvement in overall running time (including both pre-processing and the inference for the remaining unlabeled variables) and relative energy change.⁵ The baseline is expansion moves with no pre-processing. Let T_i^{ALG} and E_i^{ALG} be the running time and energy for algorithm ALG on the i -th instance. We define the *speedup* as $T_i^{\alpha\text{-EXP}}/T_i^{\text{ALG}}$ and *energy change* $(E_i^{\text{ALG}} - E_i^{\alpha\text{-EXP}})/E_i^{\alpha\text{-EXP}}$ for each instance, and then report the average speedup and energy change for the whole dataset.

We also report the percentage of labeled variables during the pre-processing. Since we view the decision problem (whether a given partial labeling is persistent) as a classification problem, we interchangeably use the term *percentage of labeled variables* and *recall* value. Getting the precision value is tricky. Since it’s a NP-hard problem so we cannot have the ground truth label for every variable. However, we apply our pre-processing technique to the binary subproblems induced from expansion moves. We know that either max-flow solves the subproblem exactly for the submodular cases or QPBO can find a sufficiently large subset of partial persistent labeling for the non-submodular cases (in our experiments, it labels almost all the variables). Therefore, we report the *precision* value of our method on the subset of the variables where we know the ground truth labeling.

⁵Our goal is efficient energy minimization, so speed and energy are the key criterion, but we also provide visual results comparison in the supplemental material.

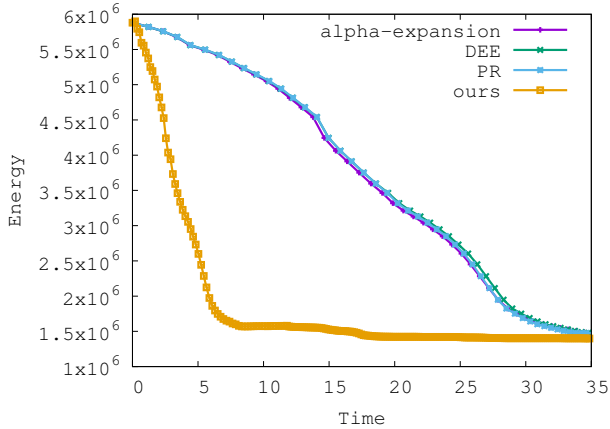


Figure 1. Energy vs time curve on instance Ted in Stereo dataset

Parameter setup and sensitivity analysis The discriminative rule in our approach has a few parameters. In order to achieve a fair comparison, we employed leave-one-out cross-validation (see, e.g. [31]) to use all but one instances in the same dataset as the validation set to choose the best parameter⁶ and test on the remaining instance. We explored all the combinations of 1) threshold $\kappa \in \{0.7, 0.8, 0.9\}$, 2) using a uniform distribution or the distribution derived from the unary term for our $q(x)$, 3) using Section 3.3 to compute LB on line 4 of Algorithm 1 or using brute force to compute $\sum_{z_{N(i)} \in \hat{\mathcal{L}}_{N(i)}(x_i=\ell)} q(z_{N(i)})$ exactly, 4) number of iterations $\tau \in \{1, 3, 5, 7, 9\}$. We run expansion moves until convergence or after 5 iterations through of the whole label set. We set our worst case bound $\epsilon = \infty$ in Section 5.2, in order to investigate how good our discriminative rule is even without the worst case guarantee. We further study the role of ϵ in Section 5.3 and supplementary material.

There is evidence that our approach achieves good performance over a wide range of parameters. We observed that cross validation picked nearly identical parameters for every instance in the same dataset. Using nearby parameters also produced good results.

We also experimented with the following fixed parameters, to avoid the expense of cross-validation: $\kappa = 0.8$, $\tau = 3$, using the uniform distribution for $q(x)$ and checking with Section 3.3. Note that this is a fairly conservative assumption, since we use the exact same parameters for very different energy functions, but still obtain good results. We achieve a 2x-12x speedup on different datasets with the energy increasing 0.1% on the worst case. In addition, we still get lower energy on 4 of the 5 challenging dataset. We defer the details of our fixed parameter experiments to the supplementary material.

⁶Based on the criterion that we choose the fastest overall running time when the false positive rate is less than 1%.

5.2. Results on benchmarks

We summarize our experimental results in Table 2. Our primary goal is to speedup MRF inference on hard problems, and there is evidence that our benchmarks are challenging. The state-of-the-art IRI method, which delivers impressive performance on the easier problems in our benchmarks, struggles with the harder problems⁷ while MPQBO runs out of memory. The only source code for Kovtun we could find is restricted to the Potts model.

Our approach achieved a significant improvement, making expansion moves 2x to 12x faster on various datasets. Our pre-processing method beats its natural competitor DEE by around 2x, and outperforms all the baseline methods. Figure 1 shows a typical energy vs. time curve. We can see our approach drives the energy curve down much faster than the other methods.

The key factor for the speedup is the percentage of labeled variables. The values of these variables are fixed during the pre-processing step, resulting in a smaller problem for max-flow/QPBO to solve. Table 2 shows our approach labels significantly more variables than DEE and PR, especially on the inpainting and denoising-sq datasets. Kovtun, MQPBO and IRI have very expensive overhead as the pre-processing step. While it is impressive that IRI labels almost every variable on the easy dataset, it is still 2x-6x slower than our proposed method. Furthermore, Kovtun, MQPBO and IRI do not perform well on our challenging datasets. When the size of the label set is large (which is common in many vision problems such as inpainting, denoising or optical flow), even IRI only proves a few variables to be persistent after spending 3x-70x as much time as our method. This demonstrates the advantage of performing pre-processing on the binary subproblem, which is consistent with the observation in [42].

Our method also performs well in terms of energy, especially on the hard benchmarks. Because we can label some variables incorrectly during pre-processing, there is a risk of producing a larger energy. However, the experimental results are reassuring: on the hard problems we actually produce slightly lower energy, while on the easier problems we can produce slightly higher energy.

While it is somewhat counter-intuitive, occasionally labeling variables incorrectly can plausibly lead to a better overall energy by getting out of a local minimum. Expansion moves can be viewed as a local search algorithm although its search space has an exponential size [1]. Therefore, a random walk going uphill occasionally may help us escape from the local minimizer, as in the Metropolis algorithm [30] or simulated annealing [16]. At one iteration of the expansion move algorithm, our method may label some

⁷In the table, time out means we do not obtain results after running for 10x the overall running time that expansion moves take.

Table 2. Experimental Results (N/A: not applicable, TO: time out, MEM: out of memory)

	Dataset	Measurement	Ours	DEE	PR	Kovtun	MQPBO	IRI
Challenging Datasets (non-Potts energy, large $ \mathcal{L} $)	Stereo	Speedup	1.78x	1.06x	1.13x	N/A	MEM	0.51x
	12–20 labels	Energy Change	-0.06%	0.00%	0.00%	N/A	MEM	-0.15%
	Trunc. L1/L2	Labeled Vars	44.76%	10.07%	18.06%	N/A	MEM	56.45%
	Inpainting	Speedup	3.40x	1.28x	1.32x	N/A	MEM	0.12x
	256 labels	Energy Change	-1.71%	0.00%	0.00%	N/A	MEM	0.00%
	Trunc. L2	Labeled Vars	74.29%	21.05%	23.75%	N/A	MEM	0.36%
	Denoising-sq	Speedup	11.83x	1.20x	1.37x	N/A	MEM	0.29x
	256 labels	Energy Change	-0.02%	0.00%	0.00%	N/A	MEM	0.00%
	L2	Labeled Vars	97.91%	16.54%	29.83%	N/A	MEM	0.39%
	Denoising-ts	Speedup	11.91x	10.53x	10.64x	N/A	MEM	0.18x
	256 labels	Energy Change	0.00%	0.00%	0.00%	N/A	MEM	-0.03%
	Trunc. L2	Labeled Vars	98.32%	95.65%	97.69%	N/A	MEM	5.85%
Easy Datasets (Potts, small $ \mathcal{L} $)	Optical Flow	Speedup	4.69x	2.63	3.40x	N/A	MEM	TO
	225 labels	Energy Change	-0.04%	0.00%	0.00%	N/A	MEM	TO
	L1	Labeled Vars	77.25%	54.34%	65.51%	N/A	MEM	TO
	Color-seg-n4	Speedup	7.02x	4.55x	6.34x	2.43x	0.37x	3.67x
	4–12 labels	Energy Change	0.00%	0.00%	0.00%	0.00%	0.00%	-0.12%
	Potts	Labeled Vars	85.74%	65.38%	77.50%	70.32%	17.27%	98.44%
	Color-seg-n8	Speedup	8.33x	5.61x	6.37x	2.33x	0.32x	1.45x
	4–12 labels	Energy Change	+0.04%	0.00%	0.00%	0.00%	0.00%	-0.10%
	Potts	Labeled Vars	90.39%	71.62%	82.05%	70.05%	17.87%	99.35%

Table 3. Precision of our method

Dataset	Stereo	Inpainting	Denoising-sq	Denoising-ts	Optical Flow	Color-seg-n4	Color-seg-n8
Precision	99.74%	96.16%	99.95%	99.79%	99.88%	99.79%	99.77%

Table 4. Precision/recall value v.s. κ (P: Precision, R: Recall)

κ		0.7	0.8	0.9	1.0
Stereo	P	90.40%	99.71%	99.41%	100.00%
	R	91.31%	56.77%	11.35%	9.26%
Inpainting	P	95.11%	99.88%	99.96%	100.00%
	R	90.51%	47.06%	25.97%	21.93%
Denoising-sq	P	99.66%	99.95%	99.95%	100.00%
	R	99.47%	97.52%	19.11%	15.15%
Denoising-ts	P	99.75%	99.95%	99.99%	100.00%
	R	98.61%	96.65%	94.99%	94.62%
Optical Flow	P	94.01%	99.50%	99.98%	100.00%
	R	99.27%	93.74%	60.85%	56.79%
Color-seg-n4	P	94.77%	99.50%	99.86%	100.00%
	R	98.52%	90.80%	77.20%	66.65%
Color-seg-n8	P	99.48%	99.76%	99.87%	100.00%
	R	92.84%	90.43%	86.92%	71.66%

variables incorrectly and solve the binary subproblem suboptimally (i.e., our pre-processing may cause the energy to increase during expansion move framework). It is plausible that this suboptimal move for the binary subproblem may also help us escape from the local minimizer. To verify this hypothesis, we experimented with a variant of our method where we reject an expansion move if it makes the energy worse. In experiments, this change led to a worse final energy. This suggests that allowing suboptimal moves is beneficial.

We believe that our method achieves competitive energy due to the very high precision, shown in Table 3. It demonstrates that our discriminative rule described in Eq. 6 is effective and powerful despite being simple and intuitive. In general, by compromising precision a little bit, we can significantly boost the recall value, as illustrated in Table 4.

In summary, our proposed method achieves a high qual-

ity trade-off between running time and energy among all the methods, particularly on challenging datasets. It runs significantly faster than its competitors and achieves an energy that is similar and sometimes even lower.

5.3. Experiments with parameters and bounds

In Section 5.2, we set the parameter $\epsilon = \infty$, and investigated how our algorithm performed without the worst case bound. We have demonstrated the proposed discriminative rule Eq. 6 itself is empirically effective. All the post-running per-instance bounds we proved in Section 4.1 are still sound, although in this variant of our method there is no worst case theoretical guarantee.

However, if we combine Eq. 6 and $\bar{\delta}_i \leq \epsilon$ as our decision rule, as described in Section 4.2, we will have the worst case bounds. We also conducted experiments with different ϵ 's. Our observation is that when $\kappa \geq 0.8$, adding the rule $\bar{\delta}_i \leq \epsilon$ has minimal effects on the speedup and energy we reported in Table 2, since our precision is already very high as shown in Table 4. However, it gives us a worst case theoretical guarantee. When $\kappa \leq 0.6$, we can observe a noticeable improvement on precision and energy change when we decrease the ϵ value with other parameters fixed. As a special case, we have a sound condition again when $\epsilon = 0$. In general, decreasing ϵ increases the running time, but the tradeoffs involved are not obvious, and we defer details to the supplementary material.

Acknowledgments: This research was supported by NSF grants IIS-1447473 and IIS-1161860 and by a Google Faculty Research Award.

References

- [1] R. K. Ahuja, Ö. Ergun, J. B. Orlin, and A. P. Punnen. A survey of very large-scale neighborhood search techniques. *Discrete Applied Mathematics*, 123(1–3):75–102, 2002. 7
- [2] K. Alahari, P. Kohli, and P. H. Torr. Dynamic hybrid algorithms for MAP inference in discrete MRFs. *TPAMI*, 32(10):1846–1857, 2010. 3, 6
- [3] E. Boros and P. L. Hammer. Pseudo-boolean optimization. *Discrete Applied Mathematics*, 123(1):155–225, 2002. 1, 2
- [4] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *TPAMI*, 23(11):1222–1239, 2001. 1, 6
- [5] D. J. Butler, J. Wulff, G. B. Stanley, and M. J. Black. A naturalistic open source movie for optical flow evaluation. In *ECCV*, pages 611–625, 2012. 6
- [6] Q. Chen and V. Koltun. Full flow: Optical flow estimation by global optimization over regular grids. In *CVPR*, pages 4706–4714, 2016. 6
- [7] M. Davis and H. Putnam. A computing procedure for quantification theory. *J. ACM*, 7(3):201–215, 1960. 1
- [8] J. Desmet, M. D. Maeyer, B. Hazes, and I. Lasters. The dead-end elimination theorem and its use in protein side-chain positioning. *Nature*, 356(6369):539–542, 1992. 1, 2, 6
- [9] P. F. Felzenszwalb and R. Zabih. Dynamic programming and graph algorithms in computer vision. *TPAMI*, 33(4):721–740, 2011. 1
- [10] A. Globerson and T. S. Jaakkola. Fixing max-product: Convergent message passing algorithms for MAP LP-relaxations. In *NIPS*, pages 553–560, 2008. 6
- [11] S. Gould, F. Amat, and D. Koller. Alphabet soup: A framework for approximate energy minimization. In *CVPR*, pages 903–910, 2009. 6
- [12] H. Ishikawa. Exact optimization for Markov Random Fields with convex priors. *TPAMI*, 25(10):1333–1336, 2003. 1
- [13] F. Kahl and P. Strandmark. Generalized roof duality. *Discrete Applied Mathematics*, 160(16):2419–2434, 2012. 2
- [14] J. H. Kappes, B. Andres, F. A. Hamprecht, C. Schnörr, S. Nowozin, D. Batra, S. Kim, T. Kroeger, B. X. Kausler, J. Lellmann, B. Savchynskyy, N. Komodakis, and C. Rother. A comparative study of modern inference techniques for discrete energy minimization problems. *IJCV*, 115(2):155–184, 2015. 1, 6
- [15] J. H. Kappes, B. Savchynskyy, and C. Schnörr. A bundle approach to efficient MAP-inference by Lagrangian relaxation. In *CVPR*, pages 1688–1695, 2012. 6
- [16] S. Kirkpatrick, C. Gelatt, and M. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983. 7
- [17] J. Kleinberg and E. Tardos. Approximation algorithms for classification problems with pairwise relationships: metric labeling and Markov Random Fields. *J. ACM*, 49(5):616–639, 2002. 6
- [18] P. Kohli, A. Shekhovtsov, C. Rother, V. Kolmogorov, and P. Torr. On partial optimality in multi-label MRFs. In *ICML*, pages 480–487, 2008. 1, 2, 6
- [19] V. Kolmogorov. Convergent tree-reweighted message passing for energy minimization. *TPAMI*, 28(10):1568–1583, 2006. 5, 6
- [20] V. Kolmogorov. Generalized roof duality and bisubmodular functions. In *NIPS*, pages 1144–1152, 2010. 2
- [21] V. Kolmogorov and C. Rother. Minimizing nonsubmodular functions with graph cuts—a review. *TPAMI*, 29(7):1274–1279, 2007. 1, 2, 6
- [22] V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *TPAMI*, 26(2):147–59, 2004. 1
- [23] N. Komodakis and G. Tziritas. Approximate labeling via graph cuts based on linear programming. *TPAMI*, 29(8):1436–1453, 2007. 5, 6
- [24] I. Kovtun. Partial optimal labeling search for a NP-hard subclass of (max,+) problems. In *Pattern Recognition*, pages 402–409, 2003. 1, 2, 6
- [25] I. Kovtun. Image segmentation based on sufficient conditions of optimality in NP-complete classes of structural labelling problem. *Ukrainian. PhD thesis. IRTC ITS National Academy of Sciences, Ukraine*, 2004. 1, 2, 6
- [26] J. Lellmann and C. Schnörr. Continuous multiclass labeling approaches and algorithms. *SIAM Journal on Imaging Sciences*, 4(4):1049–1096, 2011. 6
- [27] V. Lempitsky, C. Rother, S. Roth, and A. Blake. Fusion moves for Markov Random Field optimization. *TPAMI*, 32(8):1392–1405, 2010. 1
- [28] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, pages 416–423, 2001. 6
- [29] T. Meltzer, C. Yanover, and Y. Weiss. Globally optimal solutions for energy minimization in stereo vision using reweighted belief propagation. In *ICCV*, pages 428–435, 2005. 1
- [30] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1091, 1953. 7
- [31] K. P. Murphy. *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012. 4, 6, 7
- [32] M. L. Radhakrishnan and S. L. Su. Dead-end elimination as a heuristic for min-cut image segmentation. In *ICIP*, pages 2429–2432, 2006. 3
- [33] C. Rother, V. Kolmogorov, V. Lempitsky, and M. Szummer. Optimizing binary MRFs via extended roof duality. In *CVPR*, pages 1–8, 2007. 2
- [34] B. Savchynskyy, J. H. Kappes, P. Swoboda, and C. Schnörr. Global MAP-optimality by shrinking the combinatorial search area with convex relaxation. In *NIPS*, pages 1950–1958, 2013. 2
- [35] A. Shekhovtsov. Maximum persistency in energy minimization. In *CVPR*, pages 1162–1169, 2014. 1, 2, 6
- [36] A. Shekhovtsov, P. Swoboda, and B. Savchynskyy. Maximum persistency via iterative relaxed inference with graphical models. In *CVPR*, pages 521–529, 2015. 1, 2, 6
- [37] D. Sontag, D. K. Choe, and Y. Li. Efficiently searching for frustrated cycles in MAP inference. In *UAI*, pages 795–804, 2012. 6
- [38] D. Sontag, T. Meltzer, A. Globerson, T. Jaakkola, and Y. Weiss. Tightening LP relaxations for MAP using message passing. In *UAI*, pages 503–510, 2008. 6

- [39] P. Swoboda, B. Savchynskyy, J. H. Kappes, and C. Schnorr. Partial optimality by pruning for MAP-inference with general graphical models. In *CVPR*, pages 1170–1177, 2014. [1](#), [2](#), [6](#)
- [40] 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. *TPAMI*, 30(6):1068–1080, 2008. [1](#), [6](#)
- [41] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky. MAP estimation via agreement on (hyper)trees: Message-passing and linear-programming approaches. *IEEE Transactions on Information Theory*, 51(11), 2005. [5](#)
- [42] C. Wang and R. Zabih. Relaxation-based preprocessing techniques for Markov Random Field inference. In *CVPR*, pages 5830–5838, 2016. [1](#), [2](#), [6](#), [7](#)
- [43] Y. Weiss and W. T. Freeman. On the optimality of solutions of the max-product belief-propagation algorithm in arbitrary graphs. *IEEE Transactions on Information Theory*, 47(2):736–744, 2001. [4](#), [6](#)
- [44] T. Windheuser, H. Ishikawa, and D. Cremers. Generalized roof duality for multi-label optimization: optimal lower bounds and persistency. In *ECCV*, pages 400–413, 2012. [2](#)