
Optimality of Approximate Inference Algorithms on Stable Instances

Hunter Lang
MIT

David Sontag
MIT

Aravindan Vijayaraghavan
Northwestern University

Abstract

Approximate algorithms for structured prediction problems—such as LP relaxations and the popular α -expansion algorithm (Boykov et al. 2001)—typically far exceed their theoretical performance guarantees on real-world instances. These algorithms often find solutions that are very close to optimal. The goal of this paper is to partially explain the performance of α -expansion and an LP relaxation algorithm on MAP inference in Ferromagnetic Potts models (FPMs). Our main results give *stability* conditions under which these two algorithms provably recover the optimal MAP solution. These theoretical results complement numerous empirical observations of good performance.

1 INTRODUCTION

For many problems in machine learning, there is a large gap between the theoretical guarantees offered by the best algorithms and the empirical performance of those algorithms on real data. For instance, many MAP inference problems reduce to well-studied combinatorial optimization problems that are computationally hard in the worst-case; in practice, however, heuristic approaches often obtain solutions that far surpass their worst-case guarantees. While worst-case analysis has been the method of choice in theoretical computer science to reason about algorithms, beyond-worst-case paradigms like average-case analysis, smoothed analysis, and implicit assumptions like stability have become increasingly popular in recent years (Blum and Spencer 1995; McSherry 2001; Spielman and Teng 2004; Bilu and Linial 2010; Balcan et al. 2009). Reconciling this large gap between theory and practice is an important challenge in machine learning.

Proceedings of the 21st International Conference on Artificial Intelligence and Statistics (AISTATS) 2018, Lanzarote, Spain. PMLR: Volume 84. Copyright 2018 by the author(s).

Many tasks in modern machine learning—especially in computer vision and natural language processing—are framed and solved as structured prediction problems (Nowozin et al. 2014; Globerson et al. 2015; Tsochantzidis et al. 2005), where the local structure of an instance can be used to inform global decisions. Stereo vision is one such problem: given two input images L and R (one left, one right), the task is to output a disparity value for each pixel in the left image that tells how much that pixel moved between L and R . If two neighboring pixels have similar intensities, the output should give them similar disparities. Undirected graphical models, also known as Markov Random Fields, provide a powerful framework for performing this type of structured prediction.

Solving the MAP inference problem in a Markov Random Field (MRF) gives the maximum-probability configuration of variables (e.g., the set of pixel disparities with maximum probability) taking into account the interaction effects between nearby variables. An MRF is represented using a graph $G = (V, E)$ in which each vertex $u \in V$ represents a random variable that can take values (labels) in the discrete set $L = \{1, 2, \dots, k\}$, and edges in E represent direct dependencies between different random variables. We consider *pairwise* MRFs, where dependencies are only along edges. If we let g be a labeling that maps V to L , we can write the MAP inference task for a pairwise MRF in *energy minimization* form as follows:

$$\min_g \sum_{u \in V} c(u, g(u)) + \sum_{(u,v) \in E} \theta_{(u,v)}(g(u), g(v)), \quad (1)$$

Here we can interpret $c(u, i)$ as the “node cost” of assigning label i to vertex u and $\theta_{(u,v)}(i, j)$ as the “edge cost” of simultaneously assigning label i to u and label j to v . In stereo vision, the observed pixel intensities can be used to estimate the node and edge costs (Boykov et al. 2001). Computing the MAP assignment corresponding to (1) is known to be NP-hard for many classes of MRFs (Shimony 1994).

A well-studied special case of (1) that has been successful in practice is the *Ferromagnetic Potts Model* (FPM). Here, each edge cost function is a nonnegative

weight $w(u, v) \in \mathbb{R}_{\geq 0}$ if u and v are assigned different labels, and 0 otherwise. We can assume without loss of generality that the node costs are also nonnegative. The energy minimization problem in an FPM is:

$$\min_g Q(g) = \min_g \sum_{u \in V} c(u, g(u)) + \sum_{\substack{(u,v) \in E \\ g(u) \neq g(v)}} w(u, v), \quad (2)$$

where we define $Q(g)$ to be the objective of labeling g . The problem (2) is known in theoretical computer science as UNIFORM METRIC LABELING, and is known to be NP-hard in the worst-case (Kleinberg and Tardos 2002). While polynomial-time inference algorithms exist for MRFs with simple structure—like low tree-width or submodularity—most graphical models that arise in real-world applications do not have such simple structure. In recent years, though, much work has gone into finding tractable model classes and efficient approximation algorithms for MAP inference.

Linear programming (LP) relaxations give one such class of algorithms. These algorithms relax the MAP problem to a linear program, typically by replacing constraints of the form $x \in \{0, 1\}$ with $x \in [0, 1]$, then round the (potentially fractional) relaxed solutions back to integral ones. In fact, solutions to these linear programming relaxations often turn out to be mostly integral for instances that arise in practice on applications like stereo vision (Komodakis and Paragios 2008; Sontag et al. 2008; Werner 2010; Tarlow et al. 2011; Kappes et al. 2013). This stands in stark contrast to our theoretical understanding of linear programming relaxations on worst-case instances¹.

Introduced by Boykov et al. (2001), the α -expansion algorithm is a simple and popular combinatorial algorithm for approximate MAP inference. It works by iteratively improving an initial labeling, each time trying to find the optimal “expansion” of a label α . It is a *local search* algorithm, and it may get stuck in local energy minima. Empirically, however, α -expansion seems to avoid bad local minima. Boykov et al. (2001) apply the algorithm to stereo vision—they construct a Ferromagnetic Potts Model from the images L and R and use α -expansion to find an approximate MAP solution, which gives a disparity value for each pixel. Surprisingly, the solutions returned by α -expansion are strikingly similar to the MAP solutions, and the output of the algorithm depends very little on the initial input labeling, even though the algorithm is a 2-approximation for (2) in the worst case. This good performance has led to wide adoption of the α -expansion algorithm in practice.

¹The standard LP relaxation for UNIFORM METRIC LABELING has an integrality gap of 2 in the worst case (Kleinberg and Tardos 2002; Manokaran et al. 2008).

The near-integrality of LP relaxations and the outstanding performance of the α -expansion algorithm on real-world data lead to the following compelling question:

Question 1. *Why do heuristics for MAP inference perform so much better in practice than their worst-case theoretical guarantees suggest? Can we identify properties of real-world instances that make them tractable?*

Real world problem instances must have structure that worst-case ones do not. To reconcile this large gap between theory and practice, we study a structural property of these instances called *stability* which we think may be key to understanding their tractability.

For many real-world instances, the ground-truth corresponds to a MAP assignment (optimal solution) that “stands out”—the optimal solution is unique and robust to small changes or errors in the instance specification. The edge costs involved in the objective are often imprecise and may only be rough estimates of the similarity between endpoints. Hence we are interested in finding the optimal solution only if the instance is stable to errors or perturbations of the edge costs. Bilu and Linial (2010) introduced a formal definition of stability in the context of graph partitioning problems to capture instances with a clear “ground-truth” solution that does not change under small multiplicative perturbations to the input.

Definition 1 ((β, γ) -perturbation). Given a weight function on the edges $w : E \rightarrow \mathbb{R}_{\geq 0}$, a weight function w' is called a (β, γ) -perturbation of w iff for any $(u, v) \in E$,

$$\frac{1}{\beta}w(u, v) \leq w'(u, v) \leq \gamma w(u, v).$$

A UNIFORM METRIC LABELING instance with edge weights w is said to be (β, γ) -stable iff the optimal solution $g^* : V \rightarrow [k]$ is unique, and it remains unchanged for any (β, γ) -perturbation w' of the edge weights—that is, g^* is the unique optimal solution for the instance with edge costs w' (see Section 3.3 for a formal definition). Note that such an instance only needs to be stable to multiplicative perturbations of the edge weights and not to perturbations of the node costs.

It may sometimes be too strong to assume that the optimal solution to the perturbed instance remains completely unchanged. In practice, when the edge costs are perturbed, the optimal solution may change a bit; however, there is often a *stable region* of each instance where the MAP assignment remains optimal. We capture such instances by introducing a weaker assumption called (β, γ, S) -weak stability: roughly speaking, weakly stable instances have a region S on which good

solutions always agree with the MAP assignment, even under edge weight perturbations (see Definition 4 for a formal definition). We now present our results for stable and weakly stable instances.

Our Results. We give provable guarantees for a natural LP relaxation and the α -expansion algorithm when the input UNIFORM METRIC LABELING instance is sufficiently stable. Our first result shows that a standard LP relaxation for UNIFORM METRIC LABELING is exact on sufficiently stable instances (please see Theorem 1 for a formal statement).

Informal Theorem. *Solving the LP relaxation (3) gives the MAP solution for any $(2, 1)$ -stable instance of UNIFORM METRIC LABELING.*

Our next result gives provable guarantees in the much more general setting of weak stability: α -expansion recovers the optimal solution on the “stable region” of any UNIFORM METRIC LABELING instance (please see Theorem 2 for a formal statement).

Informal Theorem. *Given any UNIFORM METRIC LABELING instance that is $(1, 2, S)$ -weakly stable, the α -expansion algorithm recovers the optimal solution on S . If the entire instance is $(1, 2)$ -stable (i.e. $S = V$), α -expansion finds the optimal solution.*

The above theorem shows that recovering the MAP assignment in the stable region S is tractable in polynomial time. While this implies recovery of the entire MAP assignment under the standard notion of stability (i.e., $S = V$), the theorem gives recovery guarantees even when the whole instance is not stable. This gives a partial guarantee for α -expansion, an iterative hill-climbing algorithm.

While both the results require stability to multiplicative perturbations up to a factor 2, these two stability conditions ($(2, 1)$ -stability and $(1, 2)$ -stability) seem qualitatively different. We show in Section 6 that the optimality of both of these algorithms breaks down when the stability conditions of the two theorems are switched.

2 RELATED WORK

Instance stability has been studied in the context of graph partitioning problems like Max-Cut (Bilu and Linial 2010; Bilu et al. 2013; Makarychev et al. 2014) and minimum multiway cut (Makarychev et al. 2014; Angelidakis et al. 2017), clustering problems like k -means and k -median (Awasthi et al. 2012; Balcan and Liang 2016; Balcan et al. 2015; Angelidakis et al. 2017; Dutta et al. 2017), and the traveling salesman problem (Mihalák et al. 2011).

Our work is inspired by Makarychev et al. (2014) and Angelidakis et al. (2017). Makarychev et al. (2014) developed a general framework to analyze stable instances of graph partitioning problems, showing that if there exists a convex relaxation and a rounding scheme for a problem satisfying certain properties, then the convex relaxation is exact for sufficiently stable instances of the problem. They also designed a new polynomial time iterative algorithm for “weakly stable” instances of the problem, where the optimal solution can change slightly under perturbations of the weights. The amount of stability required depends on the guarantees of the rounding scheme. Makarychev et al. (2014) applied this framework to the MINIMUM MULTIWAY CUT problem (Dahlhaus et al. 1994), which is a special case of UNIFORM METRIC LABELING. They give a polynomial-time algorithm for 4-stable² instances of MINIMUM MULTIWAY CUT. Angelidakis et al. (2017) also studied MINIMUM MULTIWAY CUT and designed a better rounding scheme to give provable guarantees for $2 - 2/k$ -stable instances.

We use the same framework as Makarychev et al. (2014) and Angelidakis et al. (2017) to prove integrality of the LP relaxation (3) (Theorem 1). However, there are several new technical challenges that we need to address to prove our results, and we briefly describe them below.

Unlike the MINIMUM MULTIWAY CUT problem, there are two different costs in UNIFORM METRIC LABELING: edge weights and node costs. Our notion of stability only assumes that the optimal solution does not change under perturbations to edge weights; we make no assumptions about perturbations to node costs.

While there is a simple reduction from an instance of UNIFORM METRIC LABELING to an instance of MINIMUM MULTIWAY CUT (Boykov et al. 1998), it converts all the node costs into edge weights. Using this reduction would effectively force us to assume stability with respect to perturbations of node costs as well. Further, the reduction creates edges of very large weight, so the stability condition required becomes very stringent. To address these challenges, we first show the existence of a new rounding scheme for the standard LP relaxation that delicately trades off the loss to the LP solution on the node costs with the loss on the edge weights. By contrast, the rounding scheme of (Angelidakis et al. 2017) may incur too much loss on the node costs relative to the LP, but node costs are irrelevant for MINIMUM MULTIWAY CUT. These new rounding guarantees suffice for our theorems.

²In MINIMUM MULTIWAY CUT, the stability parameter is given as the product $\beta\gamma$, since a (β, γ) -perturbation is equivalent to a (γ, β) -perturbation in that setting.

Algorithm 1: α -expansion Algorithm

 Initialize a labeling $f : V \rightarrow L$.

 Set `continue` = True.

while `continue` **do**

 | Set `continue` = False.

 | **for** $\alpha \in L$ **do**

 | | Find the optimal α -expansion h of f .

 | | **if** $Q(h) < Q(f)$ **then**

 | | | Set $f = h$ and `continue` = True.

 | | **end**

 | **end**
end
return f .

To obtain provable optimality guarantees for the α -expansion algorithm, we relate the improvement in each step of α -expansion to the performance of a linear program on a modified instance. We carefully perturb the weights of the given instance (based on the current solution), to argue that if the current solution is *not* optimal on the stable portion of the instance, then an α -expansion step finds a new solution with smaller cost. This proves that the α -expansion algorithm recovers the optimal solution restricted to the stable region.

3 BACKGROUND

3.1 α -expansion

Solving (2) is NP-hard (Boykov et al. 2001; Kleinberg and Tardos 2002), but several efficient approximation algorithms exist. In this work, we focus on a simple combinatorial algorithm known as α -expansion, which works by iteratively improving an initial labeling with “expansion moves.”

Definition 2 (Expansion Move). Let f be an arbitrary labeling $f : V \rightarrow L$; f gives rise to a partition S_1^f, \dots, S_k^f of the vertices, where $v \in S_i^f$ if and only if $f(v) = i$. We call a labeling $g : V \rightarrow L$ an α -expansion of f if the following two conditions hold:

$$S_\alpha^f \subset S_\alpha^g; \quad S_i^g \subset S_i^f, \quad i \neq \alpha.$$

In other words, the set of vertices labeled α may grow from f to g , and all other label sets S_i may not grow.

The full procedure is described in Algorithm 1. Boykov et al. (2001) show that the optimal α -expansion move for a given labeling f and label α can be found by solving a minimum cut problem in an auxiliary graph $G_{\alpha,f}$. Finally, they show that α -expansion is a 2-approximation for (2).

Due to its simplicity, good empirical performance, and the availability of very fast implementations, Algo-

rithm 1 has seen widespread use in practice (Kolmogorov and Zabih 2004; Kolmogorov and Zabih 2002).

3.2 LP Relaxations

Linear programming (LP) relaxations are also commonly used to find approximate MAP solutions. We will make extensive use of the following LP relaxation of (2) as a tool to analyze α -expansion and as an algorithm itself:

$$\begin{aligned} \min_{\{\bar{u}\}} \quad & \sum_{u \in V} \sum_{i \in L} c(u, i) \bar{u}_i + \sum_{(u,v) \in E} w(u, v) d(u, v) \\ \text{s.t.} \quad & \sum_i \bar{u}_i = 1, \quad \forall u \in V, \forall i \in L \\ & d(u, v) = \frac{1}{2} \|\bar{u} - \bar{v}\|_1, \quad \forall (u, v) \in E \\ & \bar{u}_i \geq 0, \quad \forall u \in V, i \in L. \end{aligned} \quad (3)$$

Here \bar{u} is the length- k vector of fractional assignments at node u . Note that the second constraint can easily be linearized using edge variables. With a slight abuse of notation, we say

$$Q(\{\bar{u}\}) = \sum_{u \in V} \sum_{i \in L} c(u, i) \bar{u}_i + \sum_{(u,v) \in E} w(u, v) d(u, v).$$

Any integer labeling f is also a feasible point of (3); f corresponds to $\{\bar{u}^f\}$, where $\bar{u}_i^f = 1$ if $f(u) = i$ and 0 otherwise. In that case, the distance $d^f(u, v) = 1$ if $f(u) \neq f(v)$ and 0 otherwise, and $Q(\{\bar{u}^f\}) = Q(f)$. We say the LP relaxation is *tight* if an optimal solution to (3) is an integer solution (i.e. $\bar{u}_i \in \{0, 1\}$ for all u and i). On Ferromagnetic Potts Models, the relaxation (3) is equivalent to the *local polytope* relaxation commonly studied in MAP inference (Wainwright and Jordan 2008; Weller et al. 2016). The appendix contains a proof of that equivalence.

3.3 Stability

We now formally define stable instances of UNIFORM METRIC LABELING.

Definition 3 ($((\beta, \gamma)$ -stable). An instance of UNIFORM METRIC LABELING (G, c, w, L) with graph G , node costs c , weights w , labels L , and optimal integer solution g is called (β, γ) -stable if for any (β, γ) -perturbation w' of w , and any labeling $h \neq g$, $Q'(h) > Q'(g)$, where Q' is the objective with costs c and weights w' .

That is, g is the unique optimal solution in any (β, γ) -perturbation. Note that node costs $c(u, i)$ are not perturbed in this definition. Requiring stability under perturbations of the costs $c(u, i)$ would lead to

a stronger condition on the input instance; perturbations to w are sufficient for our theorems. As β and γ increase, the stability condition becomes increasingly strict.

The next definition captures a broader, more local version of stability, where the instance is stable with respect to a region $S \subset V$.

Definition 4 ((β, γ, S) -weakly-stable). For some set $S \subset V$, an instance (G, c, w, L) of UNIFORM METRIC LABELING with optimal solution g is said to be (β, γ, S) -weakly-stable if for any (β, γ) -perturbation w' of the weights w and any labeling $h : V \rightarrow L$,

$$h_S \neq g_S \implies Q'(h) > Q'(g). \quad (4)$$

Here h_S and g_S are the restrictions of h and g to S , and Q' is the objective with costs c and weights w' . We call S the *stable set* or *stable region* of the instance. The weak stability property says that for any (β, γ) -perturbation of the edge weights, any solution that disagrees with the optimal solution g on the stable set must have a worse objective value. Note that a (β, γ, V) -weakly-stable instance is (β, γ) -stable.

4 LP RELAXATION AND (2,1)-STABLE INSTANCES

In this section we prove that the LP relaxation (3), which as we mentioned is equivalent to the local consistency relaxation, is tight on (2,1)-stable instances. Our proof follows the framework introduced by Makarychev et al. (2014) and Angelidakis et al. (2017). We assume *for a contradiction* that the LP is fractional at some node, then use the probabilistic method to show that there must be a labeling that violates stability of the instance. To construct this violating labeling, we build a randomized rounding algorithm for (3) that provides certain probabilistic guarantees. We show that on a carefully constructed fractional input, this rounding algorithm outputs a solution that violates stability in expectation. Thus, there must be some labeling that violates stability, and therefore the optimal LP solution must take integer values at every node. The rounding algorithm defined below and the proof techniques in this section are also used in Section 5 to analyze α -expansion.

To begin, we describe the rounding procedure. This algorithm only works on inputs that are “close” to integer solutions; we will show how to construct these so-called ε -close inputs shortly.

Definition 5 (ε -close). Fix $\varepsilon < \frac{1}{2}$. A solution $\{\bar{u}\}$ to LP (3) is ε -close to an integer labeling if for each $u \in V$, there exists some j such that $\bar{u}_j \geq 1 - \varepsilon$. Because of the simplicial constraint on each \bar{u} , this index j is unique; we can therefore refer to it as $j(u)$.

Algorithm 2: Rounding Algorithm \mathcal{R}

Define P_i as the set of vertices labeled i .

Let $\varepsilon = 1/(10k)$ and $\theta = 6/(5k)$. Note $\theta > \varepsilon$.

Choose $r \in (0, \theta)$ uniformly at random.

Choose $i \in \{1, \dots, k\}$ uniformly at random.

Apply the following rule to every node $u \in V$:

If $\bar{u}_i < r$, add u to $P_{j(u)}$. Otherwise, add u to P_i .

Return the partition (P_1, \dots, P_k) .

The rounding algorithm \mathcal{R} is defined in Algorithm 2.

The following properties of \mathcal{R} will help construct a stability-violating labeling:

Lemma 1 (Rounding Guarantees). *Let h be the (random) output of Algorithm 2 on an ε -close solution $\{\bar{u}\}$. Then:*

$$\begin{aligned} \Pr[h(u) \neq j(u)] &\geq \frac{5}{6}(1 - \bar{u}_{j(u)}) \\ \Pr[h(u) = i] &\leq \frac{5}{6}\bar{u}_i, \quad \forall i \neq j(u) \\ \Pr[(u, v) \text{ not cut}] &\geq \frac{5}{6}(1 - d(u, v)) \\ \Pr[(u, v) \text{ cut}] &\leq \frac{5}{3}d(u, v), \end{aligned}$$

where $j(u)$ is the index such that $\bar{u}_{j(u)} \geq 1 - \varepsilon$.

The appendix contains a proof of these guarantees. Note that since the rounding only works on ε -close solutions, we cannot turn these properties into an approximation algorithm. We can now use \mathcal{R} to prove the main theorem of this section:

Theorem 1. *On a (2,1)-stable instance of UNIFORM METRIC LABELING with optimal integer solution g , the LP relaxation (3) is tight.*

Proof. Assume for a contradiction that the optimal LP solution $\{\bar{u}^{LP}\}$ of (3) is fractional. To construct a stability-violating labeling, we will run Algorithm 2 on a fractional labeling $\{\bar{u}\}$ constructed from $\{\bar{u}^{LP}\}$ and the optimal integer solution g . We then use Lemma 1 to show that in expectation, the output of $\mathcal{R}(\{\bar{u}\})$ must be better than the optimal integer solution in a particular (2, 1)-perturbation, which contradicts (2, 1)-stability.

Let $\{\bar{u}^g\}$ be the solution to (3) corresponding to g , and define the following ε -close solution $\{\bar{u}\}$: for every u and every i , set $\bar{u}_i = (1 - \varepsilon)\bar{u}_i^g + \varepsilon\bar{u}_i^{LP}$. Note that $\{\bar{u}\}$ is fractional and $j(u) = g(u)$ for all u .

Recall that E_g is the set of edges cut by the optimal solution g . Define the following (2, 1)-perturbation w'

of the weights w :

$$w'(u, v) = \begin{cases} w(u, v) & (u, v) \in E_g \\ \frac{1}{2}w(u, v) & (u, v) \in E \setminus E_g. \end{cases}$$

We refer to the objective with modified weights w' as Q' (that is, Q' is the objective in the instance with weights w' and costs c).

Now let $h = \mathcal{R}(\{\bar{u}\})$. To compare g and h , we will compute $\mathbb{E}[Q'(g) - Q'(h)]$, where the expectation is over the randomness of the rounding algorithm. By definition,

$$\mathbb{E}[Q'(g) - Q'(h)] = \mathbb{E}[Q'(g) - Q'(h)|h = g] \Pr(h = g) + \mathbb{E}[Q'(g) - Q'(h)|h \neq g] \Pr(h \neq g).$$

The first term of the sum above is clearly zero. Further, as $\{\bar{u}\}$ is fractional, the guarantees in Lemma 1 imply that $\Pr(h \neq g) > 0$. By (2,1)-stability of the instance, any labeling $h \neq g$ must satisfy $Q'(h) > Q'(g)$. So stability and fractionality of the LP imply $\mathbb{E}[Q'(g) - Q'(h)] < 0$.

If we compute $\mathbb{E}[Q'(g) - Q'(h)]$ and simplify using Lemma 1 and the definition of w' (see the appendix for a full derivation), we obtain:

$$\mathbb{E}[Q'(g) - Q'(h)] \geq \frac{5}{6} \left(\sum_{u \in V} c(u, g(u)) + \sum_{(u,v) \in E_g} w(u, v) - \sum_{u \in V} \sum_{i \in L} c(u, i) \bar{u}_i - \sum_{(u,v) \in E} w(u, v) d(u, v) \right)$$

The first two terms are simply $Q(g)$, and the last two are the objective $Q(\{\bar{u}\})$ of the LP solution \bar{u} . Since $\bar{u} = (1 - \varepsilon)\bar{u}^g + \varepsilon\bar{u}^{LP}$ and $Q(\{\bar{u}^{LP}\}) \leq Q(\{\bar{u}^g\})$, the convexity of the LP objective implies $Q(\{\bar{u}\}) \leq Q(\{\bar{u}^g\}) = Q(g)$. So $\mathbb{E}[Q'(g) - Q'(h)] \geq 0$. But stability of the instance and fractionality of the LP solution implied $\mathbb{E}[Q'(g) - Q'(h)] < 0$. \square

5 α -EXPANSION AND (1,2)-STABLE INSTANCES

In this section, we study the broader stability condition given by Definition 4, where the instance may only be stable with respect to some region S . We prove that for (1,2, S)-weakly-stable instances, the α -expansion algorithm recovers the optimal solution on the stable set S . In other words, given any input α -expansion is guaranteed to recover the optimal solution on the stable portion S of the instance. When $S = V$, this implies recovery of the entire optimal solution for (1,2)-stable instances.

The proof shows that as long as the current labeling maintained by α -expansion does not agree with the optimal solution on the stable set S , there must be an expansion move that decreases the objective. Boykov et al. (2001) show that α -expansion cannot terminate while there is an expansion move that decreases the objective. Hence α -expansion cannot terminate until it agrees with the optimal solution on the stable set.

To show how to construct an expansion move that decreases the objective, we will again use the probabilistic method. We actually construct the expansion move by using the LP relaxation and rounding algorithm from the previous section. Indeed, we will use the LP solution on a modified instance to construct an input $\{\bar{u}\}$ to the rounding algorithm \mathcal{R} . We show that as long as the current labeling differs from the optimal one on the stable set, there is some labeling in the support of $\mathcal{R}(\{\bar{u}\})$ that decreases the objective. The following lemma shows that every labeling in the support of \mathcal{R} is an expansion move.

Lemma 2. *Let $\{\bar{u}\}$ be an input to Algorithm 2 and let $F : V \rightarrow L$ be the integer solution to which $\{\bar{u}\}$ is ε -close. Then for all labelings h in the support of Algorithm 2, there exists an i such that h is an i -expansion of the labeling F .*

Proof. Algorithm 2 makes a random choice of label i . Then for every vertex u , it assigns either label i or label $F(u)$ to that vertex. Clearly the set of vertices labeled i by F does not decrease, and the only new label assigned is i . So the output labeling is an i -expansion of F . \square

Therefore, there must be an expansion move that decreases the objective as long as the current labeling differs from the optimal one on the stable set. We can now state the theorem.

Theorem 2. *On a (1,2, S)-weakly-stable instance (G, c, w, L) with optimal solution g , let f be the solution output by Algorithm 1. Then $f_S = g_S$. That is, α -expansion recovers the optimal solution on the stable set.³ When $S = V$, α -expansion recovers the full optimal solution.*

We remark that Algorithm 1 is known to run in polynomial time in $|V|$ and $|L|$ as long as the costs are polynomially bounded. Veksler (1999) shows that it converges in a polynomial number of iterations when the costs and weights are constant in $|V|$ and $|L|$, (or

³The proof below shows that $Q'(f) \leq Q'(g)$ (in the sense of Definition 4) is a necessary condition for local optimality of f with respect to expansion moves. Interpreted in the framework of perturbation-stability, the local optimality argument of Boykov et al. (2001) can also be shown to give Theorem 2.

are integers that are polynomially bounded in $|V|$ and $|L|$), and each iteration performs $|L|$ maximum flow computations. In practice, Algorithm 1 typically takes only 2-5 iterations to converge (Boykov et al. 2001).

Proof. We prove that as long as the labeling at the current iteration, denoted by f , satisfies $f_S \neq g_S$, there exists an expansion move of f that decreases the objective. We use Algorithm 2 as a tool to show that this expansion must exist by constructing a particular input $\{\bar{u}\}$ such that $\mathcal{R}(\{\bar{u}\})$ has better objective than f in expectation as long as $f_S \neq g_S$.

Let $\{\bar{u}^f\}$ be the solution to LP (3) corresponding to f , and recall that E_f is the set of edges cut by f . Define the following (1,2)-perturbation of the weights w :

$$w'(u, v) = \begin{cases} w(u, v) & (u, v) \in E_f \\ 2w(u, v) & (u, v) \notin E_f \end{cases}$$

Now let $\{\bar{u}^{LP}\}$ be the optimal LP solution to the instance (G, c, w', L) with these *modified* weights. We construct an ε -close input $\{\bar{u}\}$ to the rounding algorithm: $\bar{u} = (1 - \varepsilon)\bar{u}^f + \varepsilon\bar{u}^{LP}$. For each of these labelings, the distance d in the LP relaxation (3) is given by:

$$\begin{aligned} d^f(u, v) &= \frac{1}{2} \|\bar{u}^f - \bar{v}^f\|_1 = \mathbf{1}[f(u) \neq f(v)] \\ d^{LP}(u, v) &= \frac{1}{2} \|\bar{u}^{LP} - \bar{v}^{LP}\|_1 \\ d(u, v) &= \frac{1}{2} \|\bar{u} - \bar{v}\|_1. \end{aligned}$$

From the definition of $\{\bar{u}\}$ and the triangle inequality,

$$d(u, v) \leq (1 - \varepsilon)d^f(u, v) + \varepsilon d^{LP}(u, v). \quad (5)$$

Let $h = \mathcal{R}(\{\bar{u}\})$ be the random labeling output by the rounding algorithm \mathcal{R} (Algorithm 2) on input $\{\bar{u}\}$. We now show that $\mathbb{E}[Q(f) - Q(h)] > 0$. That is, in expectation the rounding algorithm produces a solution better than f in the original instance.

$$\begin{aligned} \mathbb{E}[Q(f) - Q(h)] &= \sum_{u \in V} c(u, f(u)) \Pr[h(u) \neq f(u)] \\ &\quad + \sum_{(u, v) \in E_f} w(u, v) \Pr[(u, v) \text{ not cut}] \\ &\quad - \sum_{u \in V} \sum_{i \neq f(u)} c(u, i) \Pr[h(u) = i] \\ &\quad - \sum_{(u, v) \in E \setminus E_f} w(u, v) \Pr[(u, v) \text{ cut}]. \end{aligned}$$

Applying the rounding guarantees from Lemma 1 and using the definition of w' (here we need a (1,2)-perturbation, not a (2,1)-perturbation), we obtain the

following lower bound for $\mathbb{E}[Q(f) - Q(h)]$:

$$\begin{aligned} \mathbb{E}[Q(f) - Q(h)] &\geq \frac{5}{6} \left(\sum_{u \in V} c(u, f(u)) + \sum_{(u, v) \in E_f} w'(u, v) \right. \\ &\quad \left. - \sum_{u \in V} \sum_{i \in L} c(u, i) \bar{u}_i - \sum_{(u, v) \in E} w'(u, v) d(u, v) \right) \end{aligned}$$

Writing f as $\{\bar{u}^f\}$,

$$\begin{aligned} \mathbb{E}[Q(f) - Q(h)] &\geq \frac{5}{6} \left(\sum_{u \in V} \sum_{i \in L} c(u, i) (\bar{u}_i^f - \bar{u}_i) \right. \\ &\quad \left. + \sum_{(u, v) \in E} w'(u, v) (d^f(u, v) - d(u, v)) \right) \end{aligned}$$

By (5), $d^f(u, v) - d(u, v) \geq \varepsilon (d^f(u, v) - d^{LP}(u, v))$. Additionally, $\bar{u}_i^f - \bar{u}_i = \varepsilon (\bar{u}_i^f - \bar{u}_i^{LP})$ for all u, i . Then

$$\begin{aligned} \mathbb{E}[Q(f) - Q(h)] &\geq \frac{5}{6} \varepsilon \left(\sum_{u \in V} \sum_{i \in L} c(u, i) (\bar{u}_i^f - \bar{u}_i^{LP}) \right. \\ &\quad \left. + \sum_{(u, v) \in E} w'(u, v) (d^f(u, v) - d^{LP}(u, v)) \right) \end{aligned}$$

Using the definition of Q' (the objective in the instance (G, c, w', L)),

$$\mathbb{E}[Q(f) - Q(h)] \geq \frac{5}{6} \varepsilon (Q'(f) - Q'(\{\bar{u}^{LP}\})).$$

Recall that g is the optimal solution in the original instance. Since $\{\bar{u}^{LP}\}$ is the optimal LP solution for the instance with weights w' , $Q'(\{\bar{u}^{LP}\}) \leq Q'(g)$. Combining, we obtain:

$$\mathbb{E}[Q(f) - Q(h)] \geq \frac{5}{6} \varepsilon (Q'(f) - Q'(g))$$

If $f_S \neq g_S$, by the weak stability of the instance, $Q'(f) > Q'(g)$, so in that case $\mathbb{E}[Q(f) - Q(h)] > 0$ and there must be some labeling in the support of the rounding algorithm whose objective is less than f 's. The input to the rounding algorithm was ε -close to f , so by Lemma 2, every labeling in the support is an expansion move of f . So as long as $f_S \neq g_S$, some expansion move of f decreases the objective. Boykov et al. (2001) show that α -expansion only terminates when no expansion move decreases the objective, hence $f_S = g_S$ when the algorithm terminates. \square

6 COUNTEREXAMPLES

The algorithms analyzed in Sections 4 and 5 give guarantees in two different stability settings: (2,1)-stability, for the LP, and (1,2)-stability, for α -expansion. Here we show that each algorithm does

not provably recover the optimal solution in the other stability setting. That is, the LP relaxation is not tight on (1, 2)-stable instances, and α -expansion does not always find the optimal solution on (2, 1)-stable instances. Stability was tested by checking all possible labelings in the adversarial perturbation

$$w'(u, v) = \begin{cases} \frac{1}{\beta}w(u, v) & (u, v) \in E \setminus E_g \\ \gamma w(u, v) & (u, v) \in E_g. \end{cases}$$

A proof that this is sufficient can be found in the appendix, together with the full details on how the counterexamples were generated.




Figure 1: (1, 2)-stable instance where the LP solution is fractional.

(a) Integer OPT, Obj = 8.0 (b) LP OPT, Obj = 7.5

Node	Assignment		
u	0	0	1
v	0	0	1
w	0	0	1

Node	Assignment		
u	0	0.5	0.5
v	0.5	0.5	0
w	0.5	0	0.5

Table 1: Solutions to instance in Figure 1. Entries along each row are the assignments of the corresponding node to labels 1, 2, and 3, respectively.

LP and (1, 2)-stability. Figure 1 shows a (1, 2)-stable instance with three nodes and three labels. The optimal integer solution and optimal fractional solution are shown in Table 1. The fractional solution has strictly lower objective value. Since no edges are cut ($E_g = \emptyset$), the adversarial perturbation does not change any edge weights. The optimal solution is unique, so this instance is (1, 2)-stable. Note that α -expansion is exact on this instance—no matter the starting labeling, expanding label 3 gives the optimal solution.




Figure 2: (2, 1)-stable instance where α -expansion does not find the optimal solution.

α -expansion and (2, 1)-stability. Figure 2 shows a (2, 1)-stable instance for which α -expansion does not always find the optimal solution. Table 2 shows the optimal integer solution and the solution returned when

(a) Integer OPT, Obj = 3.0 (b) Expansion Sol, Obj = 5.0

Node	Label
u	2
v	2
w	3
x	3

Node	Label
u	1
v	1
w	1
x	2

Table 2: Solutions to instance in Figure 2.

α -expansion is run with all labels initially set to 2. The LP relaxation is tight because this instance is a tree (Wainwright and Jordan 2008). Suppose α -expansion starts with every node assigned the label 2. In the first iteration, it finds the optimal expansion of label 1, and assigns label 1 to u, v , and w , but leaves x labeled 2. After this move, no expansions can decrease the objective, so the algorithm terminates. The instance is (2, 1)-stable because in the adversarial perturbation (where edges (u, v) and (w, x) have weights 2 and 1.5, and (v, w) has weight 3), the optimal solution is still to label u, v with 2 and w, x with 3.

7 CONCLUSION

We gave conditions under which two popular algorithms for MAP inference in Ferromagnetic Potts Models are exact. The results in Section 4 provide a possible avenue for explaining why LP relaxations are often tight in practice. For weakly stable instances, the results in Section 5 provide a possible explanation for the observed phenomena that the solutions output by α -expansion are often visually indistinguishable from the optimal solution and that the output does not heavily depend on the choice of initial labeling, since we prove that the algorithm always recovers the optimal solution on the stable set S regardless of the initialization. Note that on (2, 2)-stable instances, α -expansion is exact and the LP relaxation is tight.

While the α -expansion algorithm is a local search algorithm that essentially does hill-climbing with a particular set of moves, our results show that it reaches high-quality solutions on stable instances. This implies the energy landscape of stable instances has particular properties that make MAP inference tractable, and gives many directions for future work on understanding the relationship between stability and optimization.

Acknowledgments

The authors would like to thank Fredrik D. Johansson for many helpful discussions and for his feedback on drafts of this paper. This work was supported by NSF AitF awards CCF-1637585 and CCF-1723344. AV is also supported by NSF Grant No. CCF-1652491.

References

- Angelidakis, Haris, Konstantin Makarychev, and Yury Makarychev (2017). “Algorithms for stable and perturbation-resilient problems”. In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, pp. 438–451.
- Awasthi, Pranjal, Avrim Blum, and Or Sheffet (2012). “Center-based clustering under perturbation stability”. In: *Information Processing Letters* 112.1, pp. 49–54.
- Balcan, Maria-Florina and Yingyu Liang (2016). “Clustering under Perturbation Resilience”. In: Balcan, Maria-Florina, Avrim Blum, and Anupam Gupta (2009). “Approximate clustering without the approximation”. In: *Proceedings of the Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, pp. 1068–1077.
- Balcan, Maria-Florina, Nika Haghtalab, and Colin White (2015). “Symmetric and Asymmetric k -center Clustering under Stability”. In: *arXiv preprint arXiv:1505.03924*.
- Bilu, Yonatan and Nathan Linial (2010). “Are Stable Instances Easy?”. In: *Innovations in Computer Science*, pp. 332–341.
- Bilu, Yonatan et al. (2013). “On the practically interesting instances of MAXCUT”. In: *Proceedings of the Symposium on Theoretical Aspects of Computer Science*, pp. 526–537.
- Blum, Avrim and Joel Spencer (1995). “Coloring random and semi-random k -colorable graphs”. In: *J. Algorithms* 19 (2), pp. 204–234.
- Boykov, Y., O. Veksler, and R. Zabih (1998). “Markov random fields with efficient approximations”. In: *Proceedings. 1998 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No.98CB36231)*, pp. 648–655.
- (2001). “Fast approximate energy minimization via graph cuts”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23.11, pp. 1222–1239.
- Dahlhaus, E. et al. (1994). “The Complexity of Multiterminal Cuts”. In: *SIAM Journal on Computing* 23, pp. 864–894.
- Dutta, Abhratanu, Aravindan Vijayaraghavan, and Alex Wang (2017). “Clustering Stable Instances of Euclidean k -means”. In: *Advances in Neural Information Processing Systems (to appear)*.
- Globerson, Amir et al. (2015). “How hard is inference for structured prediction?”. In: *International Conference on Machine Learning*, pp. 2181–2190.
- Kappes, Joerg et al. (2013). “A comparative study of modern inference techniques for discrete energy minimization problems”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1328–1335.
- Kleinberg, Jon and Éva Tardos (2002). “Approximation Algorithms for Classification Problems with Pairwise Relationships: Metric Labeling and Markov Random Fields”. In: *J. ACM* 49.5, pp. 616–639.
- Kolmogorov, Vladimir and Ramin Zabih (2002). “Multi-camera scene reconstruction via graph cuts”. In: *Computer Vision—ECCV 2002*, pp. 8–40.
- (2004). “What energy functions can be minimized via graph cuts?”. In: *IEEE transactions on pattern analysis and machine intelligence* 26.2, pp. 147–159.
- Komodakis, Nikos and Nikos Paragios (2008). “Beyond loose LP-relaxations: Optimizing MRFs by repairing cycles”. In: *Computer Vision—ECCV 2008*, pp. 806–820.
- Makarychev, Konstantin, Yury Makarychev, and Aravindan Vijayaraghavan (2014). “Bilu-linial stable instances of max cut and minimum multiway cut”. In: *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, pp. 890–906.
- Manokaran, Rajsekar et al. (2008). “SDP gaps and UGC hardness for multiway cut, 0-extension, and metric labeling”. In: *Proceedings of the fortieth annual ACM symposium on Theory of computing*. ACM, pp. 11–20.
- McSherry, F. (2001). “Spectral Partitioning of Random Graphs”. In: *Proceedings of the 42nd IEEE symposium on Foundations of Computer Science*. FOCS ’01. Washington, DC, USA: IEEE Computer Society, pp. 529–. ISBN: 0-7695-1390-5.
- Mihalák, Matúš et al. (2011). “On the complexity of the metric tsp under stability considerations”. In: *SOFSEM 2011: Theory and Practice of Computer Science*. Springer, pp. 382–393.
- Nowozin, Sebastian et al. (2014). *Advanced Structured Prediction*. MIT Press.
- Shimony, Solomon Eyal (1994). “Finding MAPs for belief networks is NP-hard”. In: *Artificial Intelligence* 68.2, pp. 399–410.
- Sontag, David et al. (2008). “Tightening LP Relaxations for MAP using Message-Passing”. In: *24th Conference on Uncertainty in Artificial Intelligence*. AUAI Press, pp. 503–510.
- Spielman, Daniel A. and Shang-Hua Teng (2004). “Smoothed Analysis of Algorithms: Why the Simplex Algorithm Usually Takes Polynomial Time”. In: *J. ACM* 51.3, pp. 385–463.
- Tarlow, Daniel et al. (2011). “Dynamic tree block coordinate ascent”. In: *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 113–120.
- Tsochantaridis, Ioannis et al. (2005). “Large margin methods for structured and interdependent output

- variables”. In: *Journal of machine learning research* 6.Sep, pp. 1453–1484.
- Veksler, Olga (1999). “Efficient graph-based energy minimization methods in computer vision”. In:
- Wainwright, Martin J and Michael I Jordan (2008). “Graphical models, exponential families, and variational inference”. In: *Foundations and Trends® in Machine Learning* 1.1–2, pp. 1–305.
- Weller, Adrian, Mark Rowland, and David Sontag (2016). “Tightness of LP relaxations for almost balanced models”. In: *Artificial Intelligence and Statistics*, pp. 47–55.
- Werner, Tomas (2010). “Revisiting the linear programming relaxation approach to Gibbs energy minimization and weighted constraint satisfaction”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32.8, pp. 1474–1488.