# Efficient Restoration of Multicolor Images with Independent Noise

Yuri Boykov Olga Veksler Ramin Zabih*

Cornell University, USA

Keywords: Bayesian image restoration; Markov random fields; max flow-min cut

## Abstract

*We consider the problem of maximum* a posteriori *(MAP) restoration of multicolor images where each pixel has been degraded by independent arbitrary noise. We assume that the prior distribution is given by a Markov random field with only pairwise site interactions. Two classes of site interactions are considered: two-valued site interactions, which form a generalized Potts model; and linear site interactions. We give efficient algorithms based on graph cuts for both classes. The MAP estimate for a generalized Potts model can be computed by solving a multiway minimum cut problem on a graph. While this graph problem is computationally intractable, there are fast algorithms for computing provably good approximations. The MAP estimate with linear site interactions can be computed exactly by solving a minimum cut problem on a graph. This can be performed in nearly linear time.*

# 1 Introduction

Many problems in image analysis can be formulated in a Bayesian framework, as proposed by, e.g., Geman and Geman (1984) and Besag (1986) . For image restoration, the true image is viewed as the realization of a Markov random field (MRF), whose prior distribution captures the spatial smoothness of an arbitrary scene. This true image is then degraded by noise to produce the observed image. The image restoration task is to compute the maximum *a posteriori* (MAP) estimate of the true image. The major difficulty with this approach is the extremely large computational cost, since it requires solving a global optimization problem in a space with very high dimension.

Several authors have given efficient methods for MAP estimation based on graph theory, using minimum cuts (or equivalently, maximum flows). Greig *et al.* (1989) showed that for a degraded two-color scene, the binary MAP estimate can be computed exactly via graph cuts. Ferrari *et al.* have addressed the multicolor case, although under restricted noise models.

* *Address for correspondence:* Computer Science Department, Cornell University, Ithaca, NY 14853, USA, rdz@cs.cornell.edu

Ferrari *et al.* (1995) gave a method for approximating the MAP estimate for multicolored images with a Potts model prior, using a Bernoulli field noise model. With this noise model, the true intensity is observed with probability $\epsilon$, and with probability $1 - \epsilon$ a different uniformly chosen intensity is observed. Ferrari *et al.* (1997) gave an exact solution for the MAP estimate with a prior that is close to the Potts model, under a noise model which acts independently on each bit of a pixel's intensity. With this noise model (which is common in communications theory), if the true intensity at a pixel is 128, the observed intensity is as likely to be 129 as 0.

We address the multicolor restoration problem under arbitrary noise that acts independently at each pixel, using two different smoothness priors. The first smoothness prior we investigate is a generalized Potts model; under this prior, the MAP estimate can be computed by solving a multiway minimum cut problem on a graph. This graph problem was shown by Dahlhaus *et al.* (1994) to be NP-complete, so it very likely requires exponential time. However, there are fast methods for computing a provably good approximation. The second smoothness prior involves linear site interactions; here, the exact MAP estimate can be computed by solving a standard two terminal minimum cut problem, which can be done in nearly linear time.

This paper begins with a brief review of Bayesian image restoration and MRF's. In section 3 we present a generalized Potts model, show that the MAP estimate can be obtained by finding a minimum multiway cut on a graph, and propose an efficient approximating algorithm. In section 4 the linear site interaction model is described, and a method is given to efficiently compute the MAP estimate by finding a minimum cut on a graph. Finally, in section 5 some experimental results are provided.

# 2 Bayesian image restoration

Many problems in spatial statistics require estimating some spatially varying quantity from noisy measurements. Image processing is one area in which such problems arise, but there are others (see Besag (1974)) . These problems can be naturally formulated in a Bayesian framework using Markov Random Fields, which were popularized for images by Geman and Geman (1986) . In this framework, the task is to find the maximum *a posteriori* (MAP) estimate of the underlying quantity. Bayes' rule states that the posterior probability $\Pr(f|O)$ of the hypothesis $f$ given the observations $O$ is proportional to the product of the likelihood $\Pr(O|f)$ and the prior probability $\Pr(f)$. The likelihood models the sensor noise, and the prior describes preferences among different hypotheses.

A Markov Random Field has several components: a set $\mathcal{P} = \{1, \ldots, m\}$ of sites $p$, which will be pixels; a neighborhood system $\mathcal{N} = \{ \mathcal{N}_p \mid p \in \mathcal{P} \}$ where each $\mathcal{N}_p$ is a subset of pixels in $\mathcal{P}$ describing the neighbors of $p$; and a field (or set) of random variables $F = \{ F_p \mid p \in \mathcal{P} \}$. Each random variable $F_p$ takes a value $f_p$ in some set $\mathcal{L} = \{l_1, \ldots, l_k\}$ of the possible labels (i.e., intensities). Following Li (1995) a joint event $\{F_1 = f_1, \ldots, F_m = f_m\}$ is abbreviated as $F = f$ where $f = \{ f_p \mid p \in \mathcal{P} \}$ is a *configuration* of $F$, corresponding to a realization of the field. For simplicity, we will write $\Pr(F = f)$ as $\Pr(f)$ and $\Pr(F_p = f_p)$ as $\Pr(f_p)$. In

order to be an MRF, the random variables in the field $F$ must satisfy

$$\Pr(f) > 0, \quad \forall f \in \mathcal{L}^m \qquad (positivity)$$
$$\Pr(f_p | f_{\mathcal{P}-\{p\}}) = \Pr(f_p | f_{\mathcal{N}_p}), \quad \forall p \in \mathcal{P}, \qquad (Markovianity)$$

where $\mathcal{L}^m = \mathcal{L} \times \ldots \times \mathcal{L}$ and $f_S$ represents a joint event $\{ F_p = f_p \mid p \in S \}$ for any $S \subset \mathcal{P}$. The positivity condition exists for technical reasons. The Markov condition states that each random variable $F_p$ depends on other random variables in $F$ only through its neighbors in $F_{\mathcal{N}_p} = \{F_q | \, q \in \mathcal{N}_p\}$.

The key result concerning Markov Random Fields is the Hammersley-Clifford theorem, described in Besag (1986). This states that the probability of a particular configuration $\Pr(f) \propto \exp(-\sum_C V_C(f))$, where the sum is over all cliques in the neighborhood system $\mathcal{N}$. Here, $V_C$ is a *clique potential*, which describes the prior probability of a particular realization of the elements of the clique $C$. In many applications, there are only pairwise site interactions, so

$$\Pr(f) \propto \exp\left( - \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{N}_p} V_{(p,q)}(f_p, f_q) \right).$$

In general, the field $F$ is not directly observable in the experiment. We have to estimate its realized configuration $f$ based on the observation $O$, which is related to $f$ by means of the likelihood function $\Pr(O|f)$. In the context of the image restoration problem the observation $O$ is the joint event $\{I_p = i_p\}$ over all $p \in \mathcal{P}$ where $I_p$ denotes the observable intensity at pixel $p$ and $i_p$ is its particular realization. If $F_p$ denotes the true intensity at $p$ then assuming i.i.d. noise

$$\Pr(O|f) = \prod_{p \in \mathcal{P}} g(i_p, f_p)$$

where $g(i_p, f_p) = \Pr(I_p = i_p | F_p = f_p)$ represents the sensor noise model.

We wish to obtain the configuration $f \in \mathcal{L}^m$ that maximizes the posterior probability $\Pr(f|O)$. Bayes' law tells us that $\Pr(f|O) \propto \Pr(O|f)\Pr(f)$. It follows that our MAP estimate $f$ should minimize the energy function

$$E(f) = \sum_{p \in \mathcal{P}} \sum_{q \in \mathcal{N}_p} V_{(p,q)}(f_p, f_q) - \sum_{p \in \mathcal{P}} \ln\left(g(i_p, f_p)\right).$$

In sections 3 and 4 we will address two different choices of the clique potential $V_{(p,q)}$.

# 3 The generalized Potts model

In the generalized Potts model, the pairwise clique potential is based on the unit impulse function $\delta(\cdot)$. More precisely,

$$V_{(p,q)}(f_p, f_q) = u_{\{p,q\}}(1 - \delta(f_p - f_q)).$$

Note that $\{p, q\}$ is a set, not a tuple. This implies that $V_{(p,q)}(f_p, f_q) = V_{(q,p)}(f_q, f_p)$, that is, the corresponding MRF is isotropic. Since the value of $u_{\{p,q\}}$ depends on the relative

position of the clique $\{p,q\}$ in $\mathcal{P}$ then, in general, our MRF is nonhomogeneous. If the coefficient $u_{\{p,q\}} = const \geq 0$ does not depend on $\{p,q\}$, the clique potential $V_{(p,q)}$ results in the homogeneous model introduced by Potts (1952) .

The prior probability for a generalized Potts model is thus

$$\Pr(f) \propto \exp\left(-\sum_{\{p,q\}\in\mathcal{E}_{\mathcal{N}}} 2u_{\{p,q\}}(1-\delta(f_p - f_q))\right)$$

where $\mathcal{E}_{\mathcal{N}}$ is the set of distinct $\{p,q\}$ such that $q \in \mathcal{N}_p$. Each term in the summation above equals $2u_{\{p,q\}}$ if $p$ and $q$ have different labels ($f_p \neq f_q$) and zero otherwise. The coefficient $u_{\{p,q\}}$ can be interpreted as a cost of a "discontinuity" between $p$ and $q$, that is, the penalty for assigning different labels to neighboring pixels $p$ and $q$. The sum in the exponent above is proportional to the total cost of discontinuities in $f$. Thus, the prior probability $\Pr(f)$ is larger for configurations $f$ with fewer discontinuities.

The posterior energy function for a generalized Potts model is

$$E(f) \quad = \quad \sum_{\{p,q\}\in\mathcal{E}_{\mathcal{N}}} 2u_{\{p,q\}}(1-\delta(f_p - f_q)) \quad - \quad \sum_{p\in\mathcal{P}} \ln\left(g(i_p, f_p)\right). \tag{1}$$

The MAP estimate $f$ minimizes $E(f)$. Thus, it should both have a small number of discontinuities and agree with the observed data.

We now show that minimizing the energy function $E(f)$ in (1) over $f \in \mathcal{L}^m$ is equivalent to solving a multiway cut problem on a certain graph. We first give another formulation of the posterior energy minimization problem that is equivalent to (1). This formulation, shown in equation (2), reduces the search space for $f$ and simplifies our transition to the graph problem. Then we construct a particular graph, and prove that solving the multiway cut problem on this graph is equivalent to minimizing the energy function of equation (2).

## 3.1 Reformulating the energy function

We want to find $f^* \in \mathcal{L}^m$ that minimizes $E(f)$ in (1). It is straightforward to reduce the search space for $f^*$. Assuming $E(f^*)$ is finite, we can always find some constant $K(p)$ for each pixel $p$ satisfying

$$-\ln(g(i_p, f_p^*)) \quad < \quad K(p).$$

For example, if no better argument is available we can always take $K(p) = K = E(f)$ where $f$ is any fixed configuration of $F$ such that $E(f)$ is finite.

For a given collection of constants $K(p)$ we define

$$\mathcal{L}_p \;=\; \{l \in \mathcal{L} : \;\; -\ln(g(i_p, l)) < K(p)\}$$

for each pixel $p$ in $\mathcal{P}$. Each $\mathcal{L}_p$ prunes out a set of labels which cannot be assigned to $p$ in the optimal solution. For example, if we take $K(p) = E(f)$ as suggested above, then for $l \notin \mathcal{L}_p$ a single sensor noise term $-\ln(g(i_p, l))$ in (1) will exceed the total value of the posterior energy function $E(f)$ at some configuration $f$. Since $f_p^* \in \mathcal{L}_p$ then each $\mathcal{L}_p$ is a nonempty set. Define also $\bar{\mathcal{L}} = \mathcal{L}_1 \times \ldots \times \mathcal{L}_m$. Since $f^* \in \bar{\mathcal{L}}$, our search can be restricted to the set $\bar{\mathcal{L}}$.

It is possible to rewrite $-\ln(g(i_p, f_p))$ as

$$\bar{K}(p) + \sum_{\substack{l \in \mathcal{L}_p \\ l \neq f_p}} (\ln(g(i_p, l)) + K(p))$$

where $\bar{K}(p)$ is some constant that does not depend on $f_p$. It follows that minimizing $E(f)$ in (1) is equivalent to minimizing

$$\bar{E}(f) \quad = \quad \sum_{\{p,q\} \in \mathcal{E}_\mathcal{N}} 2u_{\{p,q\}}(1 - \delta(f_p - f_q)) \quad + \quad \sum_{p \in \mathcal{P}} \sum_{\substack{l \in \mathcal{L}_p \\ l \neq f_p}} h(i, p, l) \tag{2}$$

where $h(i, p, l) = \ln(g(i_p, l)) + K(p)$ and the minimization takes place over $f \in \bar{\mathcal{L}}$. Note that $h(i, p, l) > 0$ for any $p \in \mathcal{P}$ and for any $l \in \mathcal{L}_p$.

## 3.2 Multiway cut formulation

Consider a graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ with non-negative edge weights, along with a set of terminal vertices $\mathcal{L} \subset \mathcal{V}$. A subset of edges $\mathcal{C} \subset \mathcal{E}$ is called a *multiway cut* if the terminals are completely separated in the *induced* graph $\mathcal{G}(\mathcal{C}) = \langle \mathcal{V}, \mathcal{E} - \mathcal{C} \rangle$. The cost of the cut $\mathcal{C}$ is denoted by $|\mathcal{C}|$ and equals the sum of its edge weights. The *multiway cut problem* is to find the minimum cost multiway cut.

We now show that the minimization problem in (2) is equivalent to a multiway cut problem. We begin by constructing $\mathcal{G}$. We take $\mathcal{V} = \mathcal{P} \cup \mathcal{L}$. This means that $\mathcal{G}$ contains two types of vertices: *p-vertices* (pixels) and *l-vertices* (labels). Note that $l$-vertices will serve as terminals for our multiway cut problem. Two $p$-vertices are connected by an edge if and only if the corresponding pixels are neighbors in $\mathcal{N}$. Therefore, the set $\mathcal{E}_\mathcal{N}$ corresponds to the set of edges between $p$-vertices. We will refer to elements of $\mathcal{E}_\mathcal{N}$ as *n-links*. Each $n$-link $\{p, q\} \in \mathcal{E}_\mathcal{N}$ is assigned a weight

$$w_{\{p,q\}} = 2u_{\{p,q\}}. \tag{3}$$

A $p$-vertex is connected by an edge to an $l$-vertex if and only if $l \in \mathcal{L}_p$. An edge $\{p, l\}$ that connects a $p$-vertex with a terminal (an $l$-vertex) will be called a *t-link* and the set of all such edges will be denoted by $\mathcal{E}_\mathcal{T}$. Each $t$-link $\{p, l\} \in \mathcal{E}_\mathcal{T}$ is assigned a weight

$$w_{\{p,l\}} = h(i, p, l) + \sum_{q \in \mathcal{N}_p} w_{\{p,q\}}. \tag{4}$$

Note that each $p$-vertex is connected to at least one terminal since $\mathcal{L}_p$ is non-empty. No edge connects terminals directly to each other. Therefore, $\mathcal{E} = \mathcal{E}_\mathcal{N} \cup \mathcal{E}_\mathcal{T}$. Figure 1 shows the general structure of the graph $\mathcal{G}$.

Since a multiway cut separates all terminals it can leave at most one $t$-link at each $p$-vertex. A multiway cut $\mathcal{C}$ is called *feasible* if each $p$-vertex is left with exactly one $t$-link. Each feasible multiway cut $\mathcal{C}$ corresponds to some configuration $f^\mathcal{C}$ in $\bar{\mathcal{L}}$ in an obvious manner: simply assign the label $l$ to all pixels $p$ which are $t$-linked to the $l$-vertex in $\mathcal{G}(\mathcal{C})$.

**Lemma 1** *A minimum cost multiway cut $\mathcal{C}$ on $\mathcal{G}$ for terminals $\mathcal{L}$ must be feasible.*

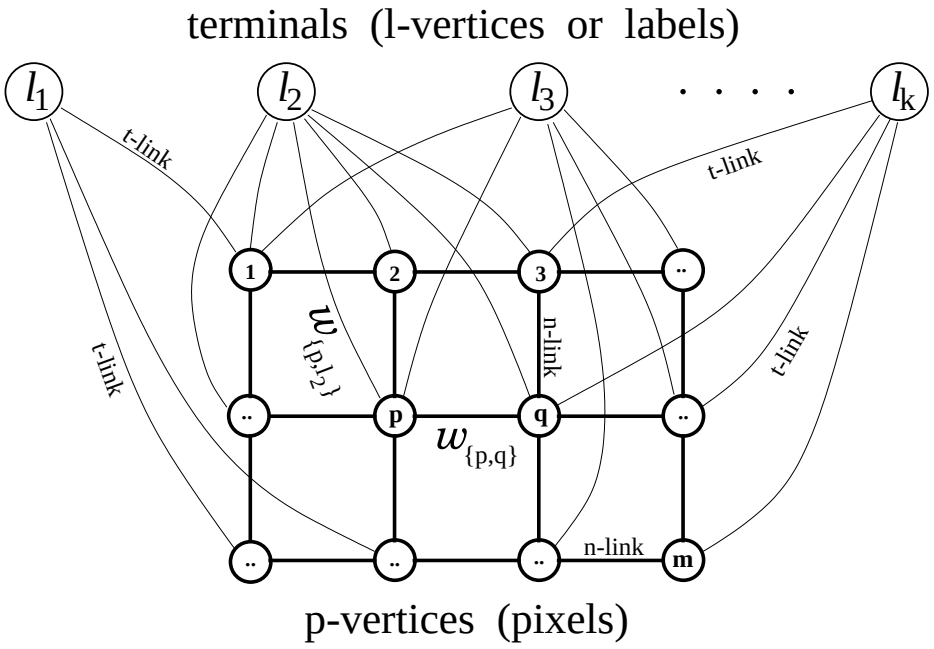


Figure 1: An example of the graph $\mathcal{G} = \langle \mathcal{V}, \mathcal{E} \rangle$ where the terminals are $\mathcal{L} = \{l_1, \ldots, l_k\}$ and $p$-vertices are elements of $\mathcal{P} = \{1, \ldots, p, q, \ldots, m\}$. Each $p$-vertex is connected to at least one terminal.

PROOF: Due to equation (4), each $t$-link $\{p, l\}$ has a weight larger then the sum of weights of all $n$-links adjacent to the $p$-vertex. If a multiway cut of minimum cost is not feasible then there exists some $p$-vertex with no $t$-link left. In such a case we will obtain a smaller cut by returning to the graph one $t$-link $\{p, l\}$ for an arbitrary $l \in \mathcal{L}_p$ and cutting all $n$-links adjacent to this $p$-vertex. ■

**Theorem 1** *If $\mathcal{C}$ is a minimum cost multiway cut on $\mathcal{G}$, then $f^{\mathcal{C}}$ minimizes $E(f)$ in (1).*

PROOF: Lemma 1 allows to concentrate on feasible multiway cuts only. Note that distinct feasible multiway cuts $\mathcal{C}1$ and $\mathcal{C}2$ can induce the same configuration $f^{\mathcal{C}1} = f^{\mathcal{C}2}$. However, among all feasible cuts corresponding to the same $f \in \bar{\mathcal{L}}$ there is a unique *irreducible* cut $\mathcal{C}$, where *irreducible* means that it does not sever $n$-links between two $p$-vertices connected to the same terminal in $\mathcal{G}(\mathcal{C})$. It follows that there is a one to one correspondence between configurations $f$ in $\bar{\mathcal{L}}$ and irreducible feasible multiway cuts on the $\mathcal{G}$.

Obviously, the minimum multiway cut should be both feasible and irreducible. To conclude the theorem it suffices to show that the cost of any irreducible feasible multiway cut $\mathcal{C}$ satisfies $|\mathcal{C}| = A + \bar{E}(f^{\mathcal{C}})$, where $A$ is the same constant for all irreducible feasible multiway cuts. Since $\mathcal{C}$ is feasible, the sum of the weights for $t$-links in $\mathcal{C}$ is equal to

$$\sum_{p \in \mathcal{P}} \sum_{\substack{l \in \mathcal{L}_p \\ l \neq f^{\mathcal{C}}_p}} w_{\{p,l\}}.$$

Since $\mathcal{C}$ is irreducible, the sum of weights for the $n$-links in the cut is equal to

$$\sum_{\{p,q\}\in\mathcal{E}_{\mathcal{N}}} w_{\{p,q\}}(1-\delta(f_p^{\mathcal{C}}-f_q^{\mathcal{C}})).$$

The theorem now follows from (3) and (4). ■

### 3.3 Multiway cut minimization

While the general multiway minimum cut problem is NP-complete, there are provably good approximations with near linear running time (see Dahlhaus *et al.*, 1992 ), and this is an area of active research. Approximating cuts, however, should be used carefully. If $\mathcal{C}$ approximates the minimum multiway cut on $\mathcal{G}$ within some known bounds, the value of $E(f^{\mathcal{C}})$ might not be within the same bounds with respect to the exact minimum of the posterior energy in (1). For example, cuts produced by the Dahlhaus algorithm are not guaranteed to be even feasible. Here we describe a method that greedily reduces the cost of multiway cuts on $\mathcal{G}$. Our algorithm generates a cut $\mathcal{C}$ such that $f^{\mathcal{C}}$ is a local minimum of the posterior energy function (1) in a certain strong sense described below.

Informally speaking, each cut separates pixels into distinct groups where pixels of each group $\mathcal{V}_l$ are assigned the same label $l$. Our algorithm considers only irreducible feasible cuts on $\mathcal{G}$. Any such cut can be uniquely represented by a feasible partition $\mathcal{P}_{\mathcal{V}} = \{\, \mathcal{V}_l \mid l \in \mathcal{L} \,\}$ of the set $\mathcal{V}$ where $l \in \mathcal{V}_l$ and $p \in \mathcal{V}_l$ implies $l \in \mathcal{L}_p$. An irreducible feasible cut $\mathcal{C}$ corresponds to $\mathcal{P}_{\mathcal{V}}$ where each $\mathcal{V}_l$ contains $l$ and all pixels $p \in \mathcal{V}$ such that $f_p^{\mathcal{C}} = l$. As an initial solution we can take any irreducible feasible cut. For example, consider a cut where $\mathcal{V}_l = \{l\} \cup \{\, p \in \mathcal{V} \mid l = \min \mathcal{L}_p \,\}$.

At each iteration we consider a fixed pair of distinct labels $\{l, \lambda\} \subset \mathcal{L}$. The basic operation is to improve the current cut $\mathcal{C}$, that is the current feasible partition $\mathcal{P}_{\mathcal{V}}$, by reallocating pixels in $\mathcal{V}_l \cup \mathcal{V}_\lambda$ between the terminals $l$ and $\lambda$. More specifically, we solve a standard two terminal min cut problem on a graph $\mathcal{G}_{\{l,\lambda\}} = \langle \mathcal{V}_{\{l,\lambda\}}, \mathcal{E}_{\{l,\lambda\}} \rangle$, where $\mathcal{V}_{\{l,\lambda\}} = \mathcal{V}_l \cup \mathcal{V}_\lambda$ and $\mathcal{E}_{\{l,\lambda\}}$ includes all edges in $\mathcal{E}$ that connect vertices in $\mathcal{V}_{\{l,\lambda\}}$. The optimal cut on $\mathcal{G}_{\{l,\lambda\}}$ divides the pixels in $\mathcal{V}_{\{l,\lambda\}}$ between the terminals $l$ and $\lambda$ and, thus, generates the new sets $\mathcal{V}_l'$ and $\mathcal{V}_\lambda'$. This yields a new feasible partition $\mathcal{P}_{\mathcal{V}}'$ corresponding to an irreducible feasible cut $\mathcal{C}'$ such that $|\mathcal{C}'| \leq |\mathcal{C}|$. If the inequality is strict, we call the iteration *successful* and accept the new cut $\mathcal{C}'$. If not, we reject the new partition and stick to the old cut $\mathcal{C}$.

At each iteration we take a new pair of terminals until all distinct pairs were considered. Then, we start a new cycle of iterations and consider the pairs of terminals all over again. The algorithm stops when no successful iterations were made in a cycle. The obtained multiway cut $\mathcal{C}$ yields $f^{\mathcal{C}}$ with the following property: the value of the energy function $E(f^{\mathcal{C}})$ cannot be decreased by switching *any* subset of pixels with one common label $l$ to any other common label $\lambda$.[1] This means that $f^{\mathcal{C}}$ achieves a local minimum of $E$ in a richer "move space" than the obvious one where a move changes the label of a single pixel. We are currently developing a more sophisticated algorithm which achieves an even stronger local minimum, where the energy function cannot be decreased by switching any set of pixels to a common label $\lambda$.

[1] If such a decrease in the energy function was possible then the algorithm would find it by successfully reallocating pixels in $\mathcal{V}_l \cup \mathcal{V}_\lambda$ between the terminals $l$ and $\lambda$.

Each cycle of the algorithm is quadratic in the number of labels and has the same effectively linear time complexity in the number of nodes as a standard min cut algorithm. If all edge weights are integers or rational numbers, the algorithm is guaranteed to terminate in finite number of cycles. In general, we do not have any bounds on the number of cycles it takes to complete the algorithm. Nevertheless, in the image restorations applications we considered the algorithm stops after three or four cycles. Moreover, most of the improvements are made at the first cycle. In the remaining cycles the number of pixels where the intensity is changed becomes increasingly insignificant. In section 5 we discuss the results in more details.

# 4 Linear site interactions

We now consider linear site interactions, where

$$V_{(p,q)}(f_p, f_q) = u_{\{p,q\}}|f_p - f_q|.$$

This yields the prior probability

$$\Pr(f) \propto \exp\left(-\sum_{\{p,q\}\in\mathcal{E}_{\mathcal{N}}} 2u_{\{p,q\}}|f_p - f_q|\right),$$

and the posterior energy function to be minimized

$$\tilde{E}(f) \quad = \quad \sum_{\{p,q\}\in\mathcal{E}_{\mathcal{N}}} 2u_{\{p,q\}}|f_p - f_q| \quad - \quad \sum_{p\in\mathcal{P}} \ln\left(g(i_p, f_p)\right). \tag{5}$$

The MAP estimate, i.e. the configuration $f$ that minimizes $\tilde{E}(f)$, can be computed by solving a standard two terminal minimum cut problem on a graph.

## 4.1 Two terminal cut formulation

To minimize $\tilde{E}(f)$ we apply our graph techniques of section 3.2. Consider a graph $\tilde{\mathcal{G}}$ defined as follows. There are two terminals: the source $R$ and the sink $S$. For each pixel $p$ we create a set of vertices $p_1, \ldots, p_{k-1}$. We connect them by $t$-links $\{t_1^p, \ldots, t_k^p\}$ where $t_1^p = \{R, p_1\}$, $t_j^p = \{p_{j-1}, p_j\}$, and $t_k^p = \{p_{k-1}, S\}$. For each pair of neighboring pixels $p, q$ and for each $j \in \{1, \ldots, k-1\}$ we create an $n$-link $\{p_j, q_j\}$ with weight $w_{\{p,q\}} = 2u_{\{p,q\}}$. Each $t$-link $t_j^p$ is assigned a weight $K_p - \ln(g(i, p, l_j))$ where $K_p$ is any constant such that $K_p > (k-1)\sum_{q\in\mathcal{N}_p} w_{\{p,q\}}$. The structure of a subgraph of $\tilde{\mathcal{G}}$ corresponding to a pair of neighboring pixels $p$ and $q$ is shown in figure 2.

A cut on the graph $\tilde{\mathcal{G}}$ will break at least one $t$-link for each pixel; we call a cut *feasible* if it breaks exactly one $t$-link for each pixel. Each feasible cut $\mathcal{C}$ corresponds to some configuration $f^{\mathcal{C}}$ where for each pixel $p$ we take $f_p^{\mathcal{C}} = l_j$ if the $t$-link $t_j^p$ is cut by $\mathcal{C}$.

**Lemma 2** *A minimum cut $\mathcal{C}$ on $\tilde{\mathcal{G}}$ must be feasible.*

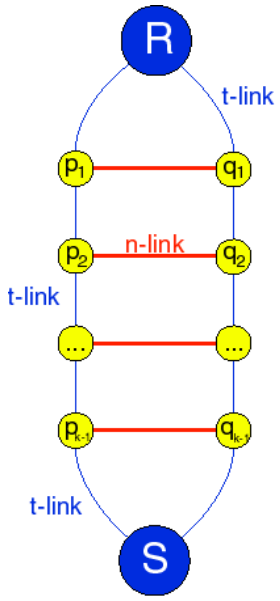


Figure 2: A subgraph of $\tilde{\mathcal{G}}$ corresponding to the pixels $p$ and $q$.

PROOF: Suppose that $t_a^p$ and $t_b^p$ are cut. Then we can find a smaller cut by restoring $t_b^p$ and breaking $n$-links $\{p_j, q_j\}$ for all $q \in \mathcal{N}_p$ and $1 \leq j \leq k-1$. The cost of the cut will decrease at least by $K_p - \ln(g(i,p,l_b)) - (k-1)\sum_{q\in\mathcal{N}_p} w_{\{p,q\}}$ which is strictly positive due to our choice of $K_p$. ■

**Theorem 2** *If $\mathcal{C}$ is a minimum cut on $\tilde{\mathcal{G}}$, then $f^{\mathcal{C}}$ minimizes $\tilde{E}(f)$ in (5).*

PROOF: We follow the same path as in the proof of Theorem 1. A cut $\mathcal{C}$ is called *irreducible* if it does not sever $n$-links between vertices connected to the same terminal in $\tilde{\mathcal{G}}(\mathcal{C})$. It is easy to show that there is a one to one correspondence between the set of all irreducible feasible cuts and the set of all configurations $f \in \mathcal{L}^m$. Since the minimum cut has to be both feasible and irreducible it remains to show that the cost of any irreducible feasible cut $\mathcal{C}$ satisfies $|\mathcal{C}| = A + \tilde{E}(f^{\mathcal{C}})$. If $\mathcal{C}$ is feasible, the cost of cutting $t$-links is $\sum_{p\in\mathcal{P}} \left(K_p - \ln(g(i,p,f_p^{\mathcal{C}}))\right)$. If $\mathcal{C}$ is also irreducible the cost of cutting $n$-links is $\sum_{\{p,q\}\in\mathcal{E}_{\mathcal{N}}} w_{\{p,q\}} |f_p^{\mathcal{C}} - f_q^{\mathcal{C}}|$. ■

Theorem 2 shows that finding an MAP estimate of the linear clique potential MRF described in the beginning of section 4 is equivalent to a standard two terminal minimum cut problem on the graph $\tilde{\mathcal{G}}$. Since $\tilde{\mathcal{G}}$ has only two terminals then the exact solution can be generated by a standard maximum flow algorithm, such as those given by Ford and Fulkerson (1962) or Goldberg and Tarjan (1988) . These algorithms have polynomial worst case complexity and have almost linear running time in many practical applications.

A graph with a similar structure was first suggested by Cox and Roy (1998) for a stereo correspondence problem. Their approach is not based on MAP-MRF estimation. The difference between $\tilde{\mathcal{G}}$ and their graph lies in the link weights. Our choice of edge weights

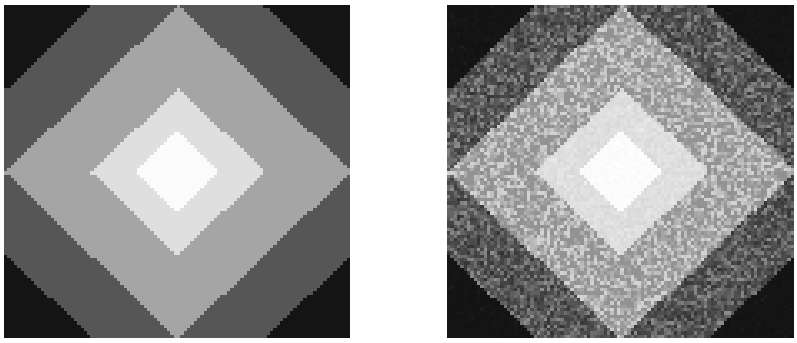

(a) The original image. The intensities of rectangles are 65, 105, 145, 185, and 225.

(b) The noisy image. The distribution of noise at each pixel is $N(m = 0, \sigma^2 = 16)$.

Figure 3: Diamond images.

guarantees the optimality property of Theorem 2. In contrast, the weights use by Roy and Cox lack theoretical justification. As a result, their algorithm does not appear to have any optimality properties.

Note that Ishikawa and Geiger (1998) describe an image segmentation technique that finds the global minimum of an energy function closely related to $\tilde{E}(f)$. Their solution, developed independently from ours, finds a minimum cut on a graph similar to $\tilde{\mathcal{G}}$ except for some details. For example, their graph is directed and has some infinite capacity links, while we employ an undirected graph.

# 5 Experimental results

In this section we discuss image restoration results in examples where an independent Gaussian noise was added to an artificial diamond image shown in Figure 3(a). The diamond image with the noise added to it is shown in Figure 3(b). For the experiments we assume that the noisy image represents the observed data and that the distribution function of noise at each pixel is known. The purpose of the experiment is to run image restoration algorithms on the observed data shown in Figure 3(b) and then to compare the restored image with the original (true) diamond image shown in Figure 3(a). The diamond image is used to test both the Potts model of section 3 and the linear clique potential model of section 4.

## 5.1 Generalized Potts model

To specify the generalized Potts models we need to choose the coefficients $u_{\{p,q\}}$. In general, the coefficients $u_{\{p,q\}}$ can be assigned different values. For example, this can be done if the image comes with some prior information about the possible locations of the object edges. In such a case, one would assign a smaller value to $u_{\{p,q\}}$ if the prior probability of an edge at $\{p, q\}$ is high and a larger value otherwise. In the experiments of this subsection we assume

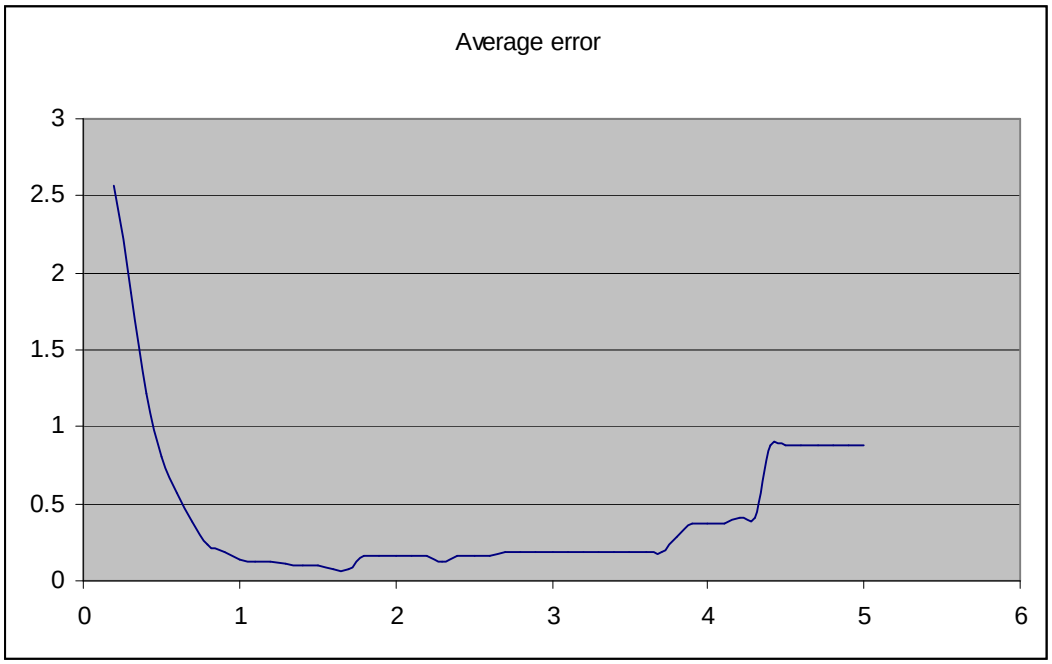


(a) Mean absolute error

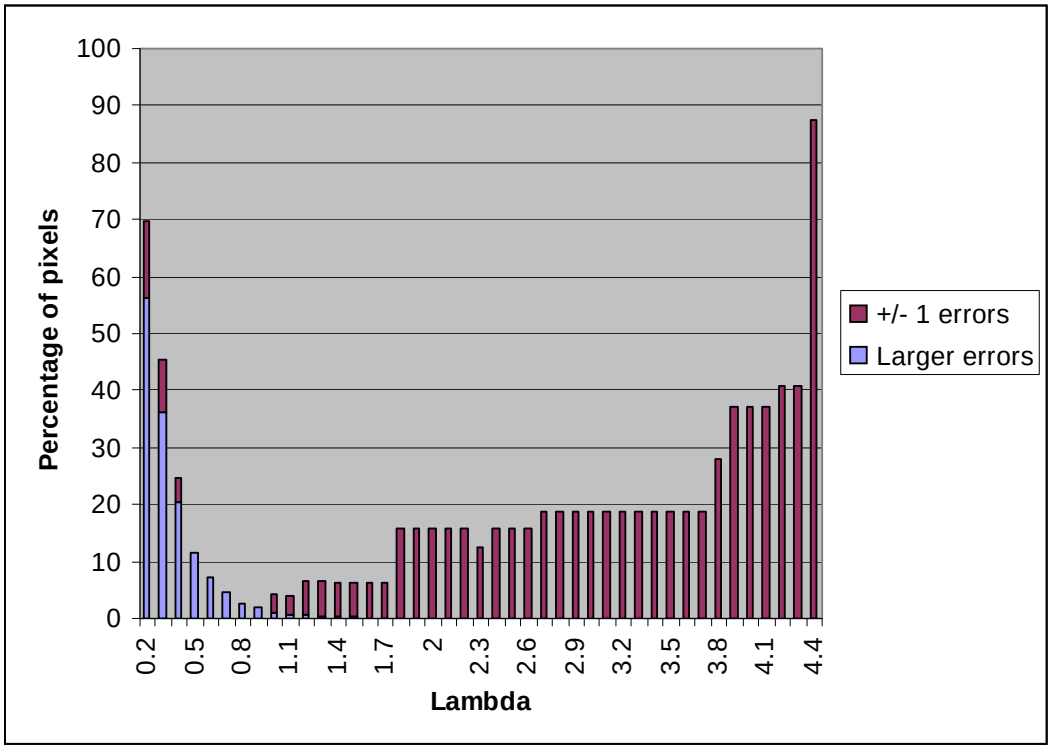


(b) Percentage of pixels with nonzero error

Figure 4: Generalized Potts model. Error analysis for various values of $\lambda$.

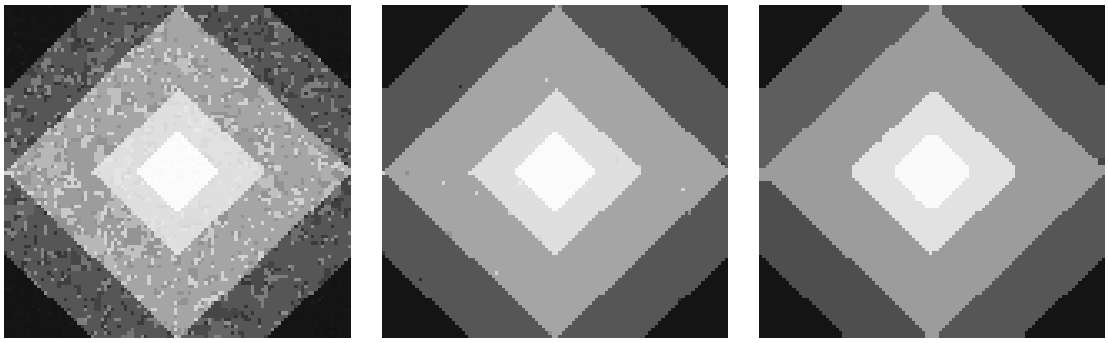

(a) $\lambda = 0.3$ (b) $\lambda = 1.4$ (c) $\lambda = 50$

Figure 5: Generalized Potts model. Restored images for various values of $\lambda$.

that no such prior information is available and that discontinuities are equally likely for all pairs of neighboring pixels. We take $u_{\{p,q\}} = \lambda$ where $\lambda$ is some nonnegative constant. We tested the algorithm of section 3.3 for various values of $\lambda$. If $\lambda$ is large then the discontinuity term of the energy function $E(f)$ in (1) become costly while the data term becomes less important. If $\lambda$ is small then the energy function $E(f)$ is more sensitive to the data term and ignores discontinuities. Informally speaking, the choice of $\lambda$ reflects our prior view on the strength of connection between the pixels of the image. Note also that $u_{\{p,q\}} = const$ implies that the specific experiments we consider in this subsection are the case of the standard homogeneous Potts model.

The images restored from the noisy data can be compared with the true diamond image. At each pixel $p$ we compute an absolute error $e_p = |i_p^t - f_p|$ where $i_p^t$ is the true intensity at $p$ and $f_p$ is the intensity at $p$ obtained by running our algorithm on the image corrupted by noise. The mean absolute error is an average of $e_p$ for all image pixels. The plot in Figure 4(a) shows the mean absolute error obtained for various values of $\lambda$. The plot in Figure 4(b) shows the number of pixels that differ from their true intensity for different values of $\lambda$. The black part of each column represents pixels with absolute errors of size one and the gray part of each column represents pixels with absolute errors larger than one.

A level of errors for very small $\lambda$'s is large since in this case the smoothness prior term of the posterior energy $E(f)$ in (1) is practically negligible. For example, for $\lambda = 0.3$ the restored image is not very different from the observed noisy data (see Figure 5(a)). If the value of $\lambda$ becomes larger then the smoothness prior term of the energy function $E(f)$ becomes more important. As a result, the restored images tend to be smoother and contain fewer discontinuities. For $\lambda = 1.4$ the restored image (see Figure 5(b)) is very close to the original diamond image. In this case only 6% of pixels on the restored image differ from their intensity in the original image and most of these errors are $\pm 1$. Note that the plots on Figure 4 show a wide range of values for $\lambda$ where the level of errors is low.

As the value of $\lambda$ grows the restored images become even smoother. The drawback, however, is that the data term of the energy function $E(f)$ becomes underrated. For example, at $\lambda = 50$ the restored image looks a bit oversmoothed (see Figure 5(c)) and the error

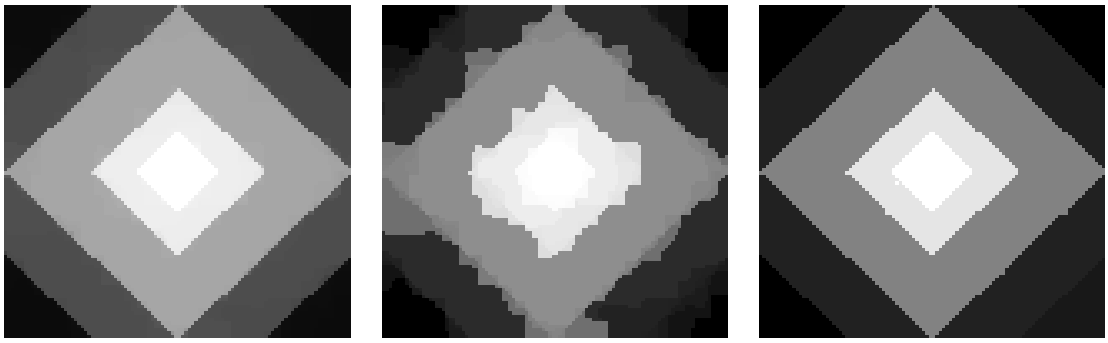

(a) Restored image.

(b) Histogram equated image. The case of constant weights $u_{\{p,q\}} = \lambda$.

(c) Histogram equated image. The case of variable weights $u_{\{p,q\}}$.

Figure 6: Linear clique potential results.

statistics noticeably deteriorates: 62% of pixels differ from the true intensity by $\pm 1$, 3% of pixels differ by $\pm 2$, 12% of pixels differ by $\pm 3$, and only 20% of pixels have the same intensity as in the original image. If we continue to increase $\lambda$ then at some point the smoothness prior term of the energy $E(f)$ would force the boundaries between the rectangles to disappear. Note, however, that such values of $\lambda$ are so large that our computer implementation starts to overflow before they are reached.

## 5.2 Linear clique potentials

Discontinuities between the neighboring pixels in the linear clique potential model are penalized differently than in the generalized Potts model. In the Potts model the cost of discontinuity between $p$ and $q$ is a fixed constant $u_{\{p,q\}}$ for any $|f_p - f_q| > 0$, whereas in the linear clique potential model the cost is $u_{\{p,q\}} \cdot |f_p - f_q|$. The constant $u_{\{p,q\}}$ can still be viewed as a coefficient reflecting the prior probability of a discontinuity between the neighboring pixels $p$ and $q$. To test the algorithm of section 4 we consider two different schemes of assigning these coefficients: in the first scheme we take $u_{\{p,q\}} = const$ for all pairs of neighboring pixels and in the second scheme $u_{\{p,q\}}$ are assigned different values according to some prior information on the likely locations of edges in the image.

In the first experiment testing the linear clique potential model we assume that $u_{\{p,q\}} = \lambda$ where $\lambda$ is some nonnegative constant. In Figure 6(a) we show the restored image obtained for $\lambda = 0.3$ which gave the best error statistics: 69% of pixels have the same intensity as in the original diamond image, 20% of pixels have $\pm 1$ errors, 6% of pixels have $\pm 2$ errors, and 3% of pixels have $\pm 3$ errors or more. At first, the restored image looks practically identical to the original image. However, the histogram equated version of the same image (Figure 6(b)) reveals a substantial oversmoothing.

In the second experiment of this subsection we assign different values to the coefficients $u_{\{p,q\}}$. We assume that some prior information on the likely location of the edges (or discontinuities) in the image is available. For simplicity, we obtain this information from a

preliminary analysis of the observed data. More specifically, we assign a large value $\lambda_l$ to $u_{\{p,q\}}$ if the difference in observed intensities $|i_p - i_q|$ is smaller than some threshold $\Delta$ and we assign a small value $\lambda_s$ to $u_{\{p,q\}}$ if the difference $|i_p - i_q|$ is larger than $\Delta$. On an intuitive level, this prior information makes it easier to cut the $n$-links between the neighboring pixels $p$ and $q$ where the prior probability of an edge is large. For $\lambda_l = 0.6$, $\lambda_s = 0.12$, and $\Delta = 16$ we obtained a restored image with the following error statistics: 93% of pixels have no error, 6% of pixels have $\pm 1$ errors, and 1% of pixels have $\pm 2$ errors or more. In this case the restored image is also indistinguishable from the true image by a naked eye but its histogram equated version (see Figure 6(c)) shows no signs of oversmoothing. Thus, the prior information about the edges helps to improve the error statistics and to avoid oversmoothing.

## 5.3 Handling smoothly changing intensity

The diamond image in Figure 3(a) has several rectangle-shaped regions where pixels have identical intensities. In this section we consider a shaded version of the diamond image where intensities smoothly change within each rectangle from left to right. The shaded diamond image is shown in Figure 7(a). The gradual change of intensity can be explained by some real life phenomenon like a change in the light condition. Therefore, the shaded diamond image provides us with a more realistic example for analysis. As before, we add an independent Gaussian noise at each pixel. The corrupted version of the image is shown in Figure 7(b).

Restoration of the shaded diamond image turns out to be a more challenging problem for the generalized Potts model algorithm. This is due to the fact that its prior is not designed to handle smooth changes in intensities. In case of constant weights $u_{\{p,q\}} = \lambda$ we obtain the best results for $\lambda = 1.0$. The restored image is shown in Figure 7(c). The Potts model algorithm tends to generate vertical stripes that follow the direction of the gradual intensity change. The error statistics is as follows: 21% of pixels have the same intensity as the original shaded diamond image, 39% of pixels have $\pm 1$ errors, 27% of pixels have $\pm 2$ errors, 13% of pixels have errors $\pm 3$ and more.

The linear clique potential model does a better job in case of the shaded diamond image. We obtained the best results for the variable weights. For simplicity, we used the same variable weights scheme as in section 5.2. The restored image is shown in Figure 7(d) and its error statistics is the following: 41% of pixels have no error, 46% of pixels have $\pm 1$ error, 10% of pixels have $\pm 2$ error, 3% of pixels have error $\pm 3$ and more. The tendency to create the stripes along the direction of the shade is significantly reduced in comparison with the generalized Potts model.

## Acknowledgments

We thank J. Kleinberg, D. Shmoys and E. Tardos for providing important references and for insightful remarks on the content of the paper. This research has been supported by DARPA under contract DAAL01-97-K-0104, monitored by ONR.

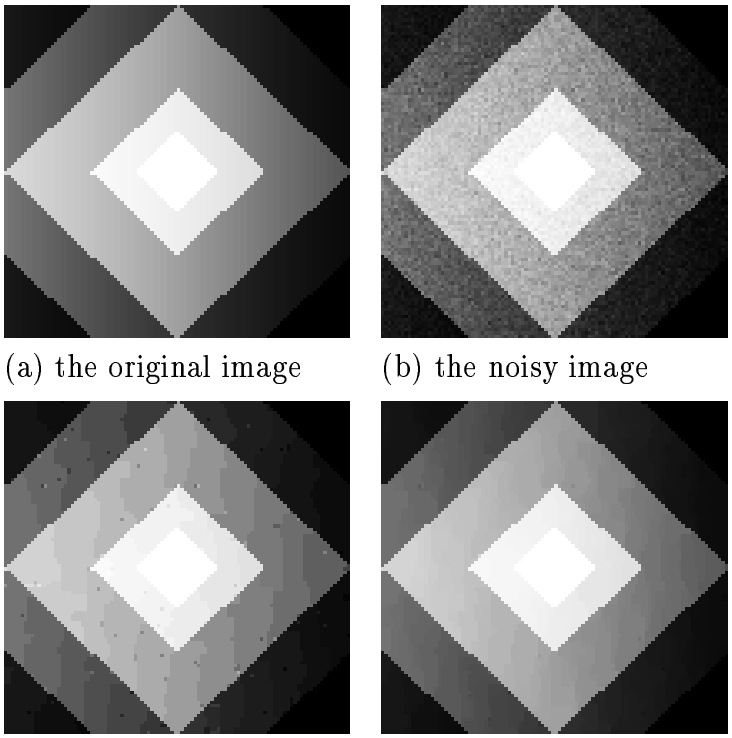

(a) the original image (b) the noisy image

(c) Potts model (d) Linear cliques

Figure 7: Shaded diamond experiments.

# References


[1] J. Besag. Spatial interaction and the statistical analysis of lattice systems. *Journal of the Royal Statistical Society, Series B*, 36:192–236, 1974.

[2] J. Besag. On the statistical analysis of dirty pictures (with discussion). *Journal of the Royal Statistical Society, Series B*, 48(3):259–302, 1986.

[3] E. Dahlhaus, D. Johnson, C. Papadimitriou, P. Seymour, and M. Yannakakis. The complexity of multiterminal cuts. *SIAM Journal on Computing*, 23(4):864–894, August 1994.

[4] E. Dahlhaus, D. S. Johnson, C.H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The complexity of multiway cuts. In *ACM Symposium on Theory of Computing*, pages 241–251, 1992.

[5] P. Ferrari, A. Frigessi, and P. de Sá. Fast approximate maximum a posteriori restoration of multicolour images. *Journal of the Royal Statistical Society, Series B*, 57(3):485–500, 1995.

[6] P. Ferrari, M. Gubitoso, and E. Neves. Restoration of multicolor images. Available from http://www.ime.usp.br/ pablo, December 1997.

[7] L. Ford and D. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.

[8] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741, 1984.

[9] A. Goldberg and R. Tarjan. A new approach to the maximum flow problem. *Journal of the Association for Computing Machinery*, 35(4):921–940, October 1988.

[10] D. Greig, B. Porteous, and A. Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society, Series B*, 51(2):271–279, 1989.

[11] H. Ishikawa and D. Geiger. Segmentation by grouping junctions. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1998.

[12] S. Li. *Markov Random Field Modeling in Computer Vision*. Springer-Verlag, 1995.

[13] R. Potts. Some generalized order-disorder transformation. *Proceedings of the Cambridge Philosophical Society*, 48(106), 1952.

[14] S. Roy and I. Cox. A maximum-flow formulation of the $n$-camera stereo correspondence problem. In *6th International Conference on Computer Vision*, 1998.