



Fast Approximate Maximum a Posteriori Restoration of Multicolour Images

Pablo A. Ferrari; Arnaldo Frigessi; Paula Gonzaga de Sa

Journal of the Royal Statistical Society. Series B (Methodological), Vol. 57, No. 3. (1995), pp. 485-500.

Stable URL:

<http://links.jstor.org/sici?sici=0035-9246%281995%2957%3A3%3C485%3AFAMAPR%3E2.0.CO%3B2-L>

Journal of the Royal Statistical Society. Series B (Methodological) is currently published by Royal Statistical Society.

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/rss.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The JSTOR Archive is a trusted digital repository providing for long-term preservation and access to leading academic journals and scholarly literature from around the world. The Archive is supported by libraries, scholarly societies, publishers, and foundations. It is an initiative of JSTOR, a not-for-profit organization with a mission to help the scholarly community take advantage of advances in technology. For more information regarding JSTOR, please contact support@jstor.org.

Fast Approximate Maximum *a Posteriori* Restoration of Multicolour Images

By PABLO A. FERRARI, ARNOLDO FRIGESSI† and PAULA GONZAGA DE SÁ

Universidade de São Paulo, Brazil

*Terza Università di Roma and
Consiglio Nazionale delle
Ricerche, Rome, Italy*

*Université Catholique de Louvain,
Louvain-la-Neuve, Belgium*

[Received September 1993. Final revision September 1994]

SUMMARY

We propose a new algorithm for the approximation of the maximum *a posteriori* (MAP) restoration of noisy images. The image restoration problem is considered in a Bayesian setting. We assume as prior distribution multicolour Markov random fields on a graph whose main restriction is the presence of only pairwise site interactions. The noise is modelled as a Bernoulli field. Computing the mode of the posterior distribution is NP complete, i.e. can (very likely) be done only in a time exponential in the number of sites of the underlying graph. Our algorithm runs in polynomial time and is based on the coding of the colours. It produces an image with the following property: either a pixel is coloured with one of the possible colours or it is left blank. In the first case we prove that this is the colour of the site in the exact MAP restoration. The quality of the approximation is then measured by the number of sites being left blank. We assess the performance of the new algorithm by numerical experiments on the simple three-colour Potts model. More rigorously, we present a probabilistic analysis of the algorithm. The results indicate that the approximation is quite often sufficiently good for the interpretation of the image.

Keywords: BAYESIAN IMAGE RESTORATION; CODING; FORD–FULKERSON'S MAX-FLOW ALGORITHM; MARKOV RANDOM FIELDS; PROBABILISTIC ANALYSIS OF ALGORITHMS; SIMULATED ANNEALING

1. INTRODUCTION

In recent years, Bayesian methods have been extensively applied in many image restoration problems. (See the survey paper by Geman (1990).) The use of Markov random fields as prior models has been particularly successful. The true unknown image is thought of as being a realization of a Markov random field, defined to capture prior information concerning the observed scene. The underlying graph or lattice structure, the number of colours and the type of interaction are all important aspects in modelling such a spatial prior. Additional assumptions on the type of degradation and noise, which deform the true image, allow us to compute the posterior distribution given the observed image. If the deformation acts locally, then this posterior measure is still a Markov random field. Inference is then guided by the choice of a loss function. If a simple 0–1 loss function is assumed, then the optimal Bayesian estimator is the mode of the posterior distribution. This is called maximum *a posteriori* (MAP) restoration. The MAP estimator does not need to be unique: often we are interested in finding one of the absolute maxima.

†Address for correspondence: Dipartimento di Matematica, Terza Università di Roma, via C. Segre 2, I-00146 Roma, Italy.

E-mail: frigessi@iac.rm.cnr.it

Clearly, computing such an MAP estimator is a global optimization problem in a highly dimensional space. For instance, if the image is modelled on a regular grid with N sites (called pixels) and each pixel may take c colours, the space of all possible images Ω has c^N elements. Enumeration quite soon becomes unfeasible, and in fact the maximization problem is generally NP complete (see Garey and Johnson (1979)). Simulated annealing (see Gidas (1991) for a recent review) is an iterative stochastic algorithm that converges weakly to a distribution concentrated on the set of global maxima Ω^* . This is true if a certain parameter, called temperature, which controls the algorithm decreases towards 0 as $1/\log t$, where t is the iteration index. Hence simulated annealing with such a logarithmic schedule is extremely slow. Therefore faster schedules have been proposed (see Aarts and Korst (1989)), but using them does not guarantee convergence to Ω^* . The estimate obtained by running the simulated annealing algorithm with these schedules is hence an approximation to the MAP restoration. No rigorous analysis of the error therein is available; however, many simulation studies have been carried out for specific distributions. Nevertheless there does not seem to be conclusive evidence on the reliability of such approximations. In Greig *et al.* (1989) evidence is given that no 'good' MAP estimate can be obtained by simulated annealing with neither logarithmic nor faster schedules in finite time, except for distributions that are relatively close to a sitewise independent field.

We stress the importance of this point. Bayesian modelling allows a precise characterization of the properties of estimators, thus permitting the user to be confident in its conclusions. However, the algorithms used will not produce such a reliable estimator but just an approximation to it, whose quality cannot be assessed. Because of its computational complexity, we must rely on an approximate MAP restoration; but we would like to quantify and maybe to determine on the image the possibly misclassified pixels. This is the aim of this work: we propose a fast method that will produce an approximate MAP estimate, with the property that if a pixel is restored, i.e. a colour assigned to it, then that colour appears in an exact MAP estimate. Some pixels in the approximate estimate will be left blank: no colour will be assigned to them. For those pixels our method is unable to discriminate. We call and code such blank pixels as 'question-marks' (?).

We immediately start with an example. In Fig. 1(b) a three-colour image degraded by Bernoulli noise is represented. Fig. 1(c) shows the output of our restoration, where question-marks are displayed. Fig. 1(a) is the true unknown noiseless image (used to generate Fig. 1(b)).

Clearly, the crucial question is how many ?s are there in our restorations? If too many the method does not give a proper answer. However, if their number is limited then the estimate is useful. In Fig. 1(c) 3% of the pixels are ?s.

We present the results of a simulation study that shows how the number of ?s varies as a function of the variance of the noise and the attraction of the random field. Although we present cases where every pixel is left undecided, in most common situations when the signal-to-noise ratio is not too small, the number of ?s is very limited (less than 10% of the N pixels). The geometry of the ?-clusters and their location is also remarkable: ?s appear on the boundary of well-restored areas and the number of pixels that contribute to a ?-cluster is very small. This brings out another important feature of our method. If it is required to restore all pixels it is possible to treat each ?-cluster independently of the others (given the

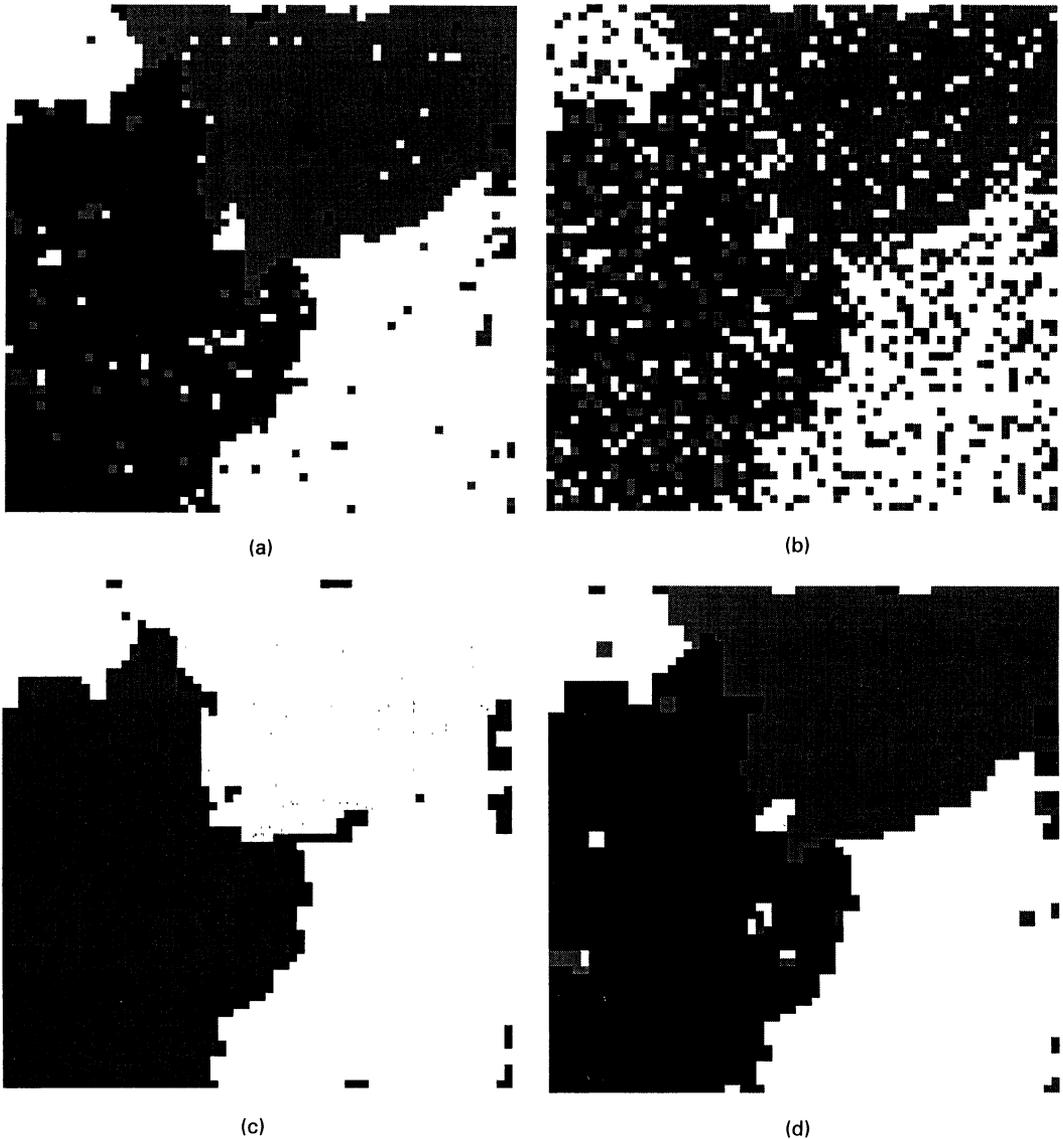


Fig. 1. (a) Original image with $\beta = 1.2$ produced by the Gibbs sampler run for 1000 raster scans and three colours; (b) image in (a) degraded with $\epsilon = 0.2$; (c) approximate MAP restoration of (b) (?s are shown in black, whereas the three colours take grey levels); (d) simulated annealing (with geometric schedule) restoration of (b)

restored boundary conditions), even in parallel, reducing the overall restoration time drastically. In particular, small clusters can be treated by full enumeration. We also note that the special location of ?s may be such that in practical situations no further restoration is required, since the ? can be resolved 'visually'. Furthermore, if large ?-areas are left, this may be considered as a sign of a very unfavourable signal-to-noise ratio and could suggest that more data must be taken to make inference more reliable.

Together with the computational experiments we also perform a rigorous probabilistic analysis of our method. In the Bayesian setting the number of ?s is considered very naturally as a random variable with respect to a probability measure P . (Precise definitions will be given later.) We prove that the expected number of ?s is CN , where C is a constant (smaller than 1) which depends on the variance of the noise and the inverse temperature of the Markov random field, which governs the strength of attraction. At a first glance CN seems quite a large number. But we should not forget the NP-completeness of the MAP restoration task. Since NP-problems are unlikely to be solved in polynomial time, we could not expect the number of ?s to be of any order smaller than $O(N^\alpha)$ for $\alpha > 0$. Unfortunately for our method we prove $\alpha = 1$. However, the constant C is smaller than 1; moreover it tends to 0 exponentially fast as the inverse temperature increases. It also decreases to 0 as the variance of the noise decreases. For the model that was used to produce Fig. 1(b), we estimated $C = 0.03$. Viewed in this light the gain is considerable: 3% pixels were left for restoration after our method had been applied.

Computing an MAP estimate is not always NP complete. In special situations the problem becomes polynomial. An important example is the Ising model with Bernoulli noise (Barahona, 1982). The Ising model is a binary ($c = 2$) Markov random field on a finite subset of \mathbb{Z}^2 with nearest neighbour interactions. Greig *et al.* (1989) reduced the MAP computation in this case to a minimum cut problem on a graph. The minimal cut can be determined very efficiently, for instance by means of the standard Ford–Fulkerson algorithm, which runs in $O(N^2)$ iterations. So in Greig *et al.* (1989) it is shown how exact MAP restorations can be obtained if the true image has exactly two colours, the prior Markov random field model features only pairwise interactions and the noise is sitewise independent (but can be signal dependent and non-symmetric). Several improvements on the original Ford–Fulkerson algorithm have been proposed in Greig *et al.* (1989) and in Jubb (1989). The latter also includes a treatment of the three-colour restoration problem: the data image is partitioned into subimages that are likely to be composed of two colours only. Then these subimages are restored and pasted together. Unfortunately these operations have an uncontrollable effect on the final estimate of the MAP restoration. The quality of the approximation cannot be quantified.

Our method, based on the results in Greig *et al.* (1989), allows us approximately to restore images with more than two colours. However, we cannot treat continuous noise, as was done in Greig *et al.* (1989). In this paper we first describe our approach for the three-colour Potts model with independently and identically distributed noise. (The MAP calculation is NP complete in this case; Barahona (1982).) Essentially we code the three colours into two colours in the three possible ways and apply the Ford–Fulkerson algorithm three times to the coded two-colour images. The three binary restorations thus obtained are then combined into a three colour plus ? image. We prove that if a colour is returned at a pixel this is the correct MAP restoration.

Section 2 includes a description of the coding and the proof. Section 3 is devoted to the probabilistic analysis of our approximation method, by estimating the expected number of ?s. We describe how the method is extended to more than three colours in Section 4. Section 5 reports on our simulations.

2. APPROXIMATE MAXIMUM A POSTERIORI RESTORATION FOR THREE-COLOUR POTTS MODELS CORRUPTED BY NOISE

2.1. Model

Let Λ be a finite subset of \mathbb{Z}^2 (e.g. a square) with N sites. Let $\mathbf{X} = (X_i)_{i=1, \dots, N}$ for $X_i \in \{0, 1, \dots, c-1\}$, $\forall i \in \Lambda$, represent the true unknown image with prior distribution $P(\mathbf{X})$ on the set of all possible images Ω . Denote the observed image by $\mathbf{Y} = (Y_i)_{i=1, \dots, N}$ and let $P(\mathbf{Y}|\mathbf{X})$ be its likelihood. The MAP estimator $\hat{\mathbf{X}} \in \Omega$ is any image that maximizes the posterior distribution $P(\mathbf{X}|\mathbf{Y}) \propto P(\mathbf{Y}|\mathbf{X})P(\mathbf{X})$:

$$P(\hat{\mathbf{X}}|\mathbf{Y}) = \max_{\mathbf{X} \in \Omega} \{P(\mathbf{X}|\mathbf{Y})\}. \quad (2.1)$$

We now specify this general setting to the three-colour Potts model observed through a memoryless symmetric channel. We take $c = 3$ and

$$P(\mathbf{X}) = \frac{1}{Z_\beta} \exp \left\{ \beta \sum_{\langle ij \rangle} \mathbf{1}_{\{X_i = X_j\}}(\mathbf{X}) \right\}, \quad \beta \geq 0, \quad (2.2)$$

where the sum is over all unordered nearest neighbouring pairs $i, j \in \Lambda$. This means that every pair is considered only once. (Boundary conditions are free.) Z_β is the normalizing constant. β is called the inverse temperature.

Next we define the noise. Conditional on \mathbf{X} , at each pixel independently the colour $Y_i \in \{0, 1, \dots, c-1\}$ is observed exactly with probability $1 - \epsilon$, whereas with probability $\epsilon/(c-1)$ it is switched to any of the remaining colours:

$$P_i(Y_i|X_i) = \begin{cases} 1 - \epsilon & \text{if } Y_i = X_i, \\ \frac{\epsilon}{c-1} & \text{otherwise.} \end{cases} \quad (2.3)$$

Hence

$$P(\mathbf{Y}|\mathbf{X}) = \prod_{i \in \Lambda} P_i(Y_i|X_i) = \frac{1}{Z_h} \exp \left\{ h \sum_{i \in \Lambda} \mathbf{1}_{\{X_i = Y_i\}}(\mathbf{Y}) \right\} \quad (2.4)$$

where

$$h = \log \left\{ \frac{(c-1)(1-\epsilon)}{\epsilon} \right\}. \quad (2.5)$$

Maximizing the posterior distribution is equivalent to maximizing the posterior energy

$$L(\mathbf{X}|\mathbf{Y}) = \beta \sum_{\langle ij \rangle} \mathbf{1}_{\{X_i = X_j\}}(\mathbf{X}) + h \sum_{i \in \Lambda} \mathbf{1}_{\{X_i = Y_i\}}(\mathbf{X}) \quad (2.6)$$

with respect to $\mathbf{X} \in \Omega$.

2.2. Coding and Decoding

Let $c = 2$. We transform the data image \mathbf{Y} into three two-colour images. This is how the coding acts. For each colour $\alpha \in \{0, 1, 2\}$ define $C^\alpha: \Omega \rightarrow Q' = \{0, 1\}^\Lambda$ sitewise as follows:

$$C^\alpha(\mathbf{X}) = (X'_i)_{i \in \Lambda}, \quad X'_i \in \{0, 1\},$$

where

$$X'_i = \begin{cases} 0 & \text{if } X_i = \alpha, \\ 1 & \text{otherwise.} \end{cases} \tag{2.7}$$

From the data image \mathbf{Y} we thus construct $C^0(\mathbf{Y})$, $C^1(\mathbf{Y})$ and $C^2(\mathbf{Y})$ in Ω' . Of course this coding does not depend on the numbering assigned to the three original colours. For each of these coded images consider the two-colour maximization problem

$$L\{\mathbf{Z}^\alpha | C^\alpha(\mathbf{Y})\} = \max_{\mathbf{Z}' \in \Omega'} [L\{\mathbf{Z}' | C^\alpha(\mathbf{Y})\}] \tag{2.8}$$

where L is defined in equation (2.6) and β and ϵ are those in the original model (2.2) and (2.3). Denote by $\mathbf{Z}^\alpha \in \Omega'$, $\alpha \in \{0, 1, 2\}$, any image realizing the maximum in equation (2.8). Computing equation (2.8) is easy by Greig *et al.* (1989) and can be performed quickly by a Ford-Fulkerson algorithm.

Of course the mapping (2.7) cannot be inverted but we can consider the set of all possible three-colour images that, once coded by conditions (2.7) to a two-colour image, are exactly equal to \mathbf{Z}^α :

$$\mathscr{H}^\alpha = \mathscr{H}^\alpha(\mathbf{Y}) = \{\mathbf{X} \in \Omega: C^\alpha(\mathbf{X}) = \mathbf{Z}^\alpha\}, \quad \alpha \in \{0, 1, 2\}. \tag{2.9}$$

It holds that, $\forall \mathbf{W} \in \mathscr{H}^\alpha$,

$$L\{C^\alpha(\mathbf{W}) | C^\alpha(\mathbf{Y})\} = L\{\mathbf{Z}^\alpha | C^\alpha(\mathbf{Y})\} = \max_{\mathbf{X} \in \Omega} [L\{C^\alpha(\mathbf{X}) | C^\alpha(\mathbf{Y})\}] \tag{2.10}$$

for $\alpha \in \{0, 1, 2\}$. If a pixel i takes colour α for an image $\mathbf{W} \in \mathscr{H}^\alpha$, then all other images in \mathscr{H}^α have pixel i with the same colour α , because $Z_i^\alpha = 0$. The images $\mathbf{W} \in \mathscr{H}^\alpha$ may differ only over sites j for which $C^\alpha(\mathbf{W})_j = 1$. We are ready to state our main result.

Theorem 1. For $\alpha \in \{0, 1, 2\}$ let $\mathbf{W} \in \mathscr{H}^\alpha$. If $W_i = \alpha$ then $\hat{X}_i = \alpha$ where

$$L(\hat{\mathbf{X}} | \mathbf{Y}) = \max_{\mathbf{X} \in \Omega} \{L(\mathbf{X} | \mathbf{Y})\},$$

i.e., if the colour α is assigned to a site $i \in \Lambda$ in the set of images \mathscr{H}^α (associated with that same colour α) then this is precisely the colour of site i in the MAP estimate of the original three-colour image.

Proof. Fix $\alpha \in \{0, 1, 2\}$. For any $\mathbf{X} \in \Omega$ let

$$\Lambda^\alpha(\mathbf{X}) = \{i \in \Lambda: X_i = \alpha\}$$

and observe that the pixel set

$$\hat{\Lambda}^\alpha = \Lambda^\alpha(\mathbf{W}), \quad \forall \mathbf{W} \in \mathscr{H}^\alpha,$$

does not depend on $\mathbf{W} \in \mathscr{H}^\alpha$. Consider now the set of images in Ω for which all pixels in $\hat{\Lambda}^\alpha$ at least take colour α

$$\mathscr{H}_+^\alpha = \{\mathbf{X} \in \Omega: \Lambda^\alpha(\mathbf{X}) \supset \hat{\Lambda}^\alpha\} \supset \mathscr{H}^\alpha.$$

The theorem is proved if we show that for any MAP estimate $\hat{\mathbf{X}} \in \Omega$, for which

$$L(\hat{\mathbf{X}}|\mathbf{Y}) \geq L(\mathbf{X}|\mathbf{Y}), \quad \forall \mathbf{X} \in \Omega, \tag{2.11}$$

it is $\hat{\mathbf{X}} \in \mathcal{S}_+^\alpha$.

Consider, for any $\mathbf{X} \in \Omega$,

$$L\{C^\alpha(\mathbf{X})|C^\alpha(\mathbf{Y})\} = L(\mathbf{X}|\mathbf{Y}) + H^\alpha(\mathbf{X}|\mathbf{Y})$$

where

$$H^\alpha(\mathbf{X}|\mathbf{Y}) = \beta \sum_{\langle ij \rangle \subset \Lambda} \mathbf{1}_{\{X_i \neq X_j, X_i \neq \alpha, X_j \neq \alpha\}} + h \sum_{i \in \Lambda} \mathbf{1}_{\{X_i \neq Y_i, X_i \neq \alpha, Y_i \neq \alpha\}}.$$

For any subset $B \subset \Lambda$ define \mathbf{X}_B as the restriction of \mathbf{X} to B . We extend the definition of L, H^α , etc. to configurations \mathbf{X}_B by restricting the sums to labels $i, j \in B$. By definition of Λ^α it follows that

$$H^\alpha(\mathbf{X}_{\Lambda^\alpha(\mathbf{X})}|\mathbf{Y}_{\Lambda^\alpha(\mathbf{X})}) = 0, \quad \forall \mathbf{X} \in \Omega. \tag{2.12}$$

From equation (2.10), for any $\mathbf{W} \in \mathcal{S}^\alpha$ we have

$$L(\mathbf{W}|\mathbf{Y}) + H^\alpha(\mathbf{W}|\mathbf{Y}) = \max_{\mathbf{X} \in \Omega} \{L(\mathbf{X}|\mathbf{Y}) + H^\alpha(\mathbf{X}|\mathbf{Y})\}$$

and in particular

$$L(\mathbf{W}|\mathbf{Y}) + H^\alpha(\mathbf{W}|\mathbf{Y}) \geq L(\tilde{\mathbf{W}}|\mathbf{Y}) + H^\alpha(\tilde{\mathbf{W}}|\mathbf{Y})$$

where we choose $\tilde{\mathbf{W}} \in \Omega$ such that $\tilde{\mathbf{W}}_{\Lambda \setminus \hat{\Lambda}^\alpha} = \mathbf{W}_{\Lambda \setminus \hat{\Lambda}^\alpha}$, i.e. any image that outside $\hat{\Lambda}^\alpha$ coincides with $\mathbf{W} \in \mathcal{S}^\alpha$, and otherwise is free. Splitting the summations in equation (2.6) (and counting boundary conditions properly only once in the $\hat{\Lambda}^\alpha$ -term), we rewrite the last inequality as

$$\begin{aligned} & \bar{L}(\mathbf{W}_{\hat{\Lambda}^\alpha}, \mathbf{W}_{\partial \hat{\Lambda}^\alpha}|\mathbf{Y}_{\hat{\Lambda}^\alpha}) + L(\mathbf{W}_{\Lambda \setminus \hat{\Lambda}^\alpha}|\mathbf{Y}_{\Lambda \setminus \hat{\Lambda}^\alpha}) + \bar{H}^\alpha(\mathbf{W}_{\hat{\Lambda}^\alpha}, \mathbf{W}_{\partial \hat{\Lambda}^\alpha}|\mathbf{Y}_{\hat{\Lambda}^\alpha}) + H^\alpha(\mathbf{W}_{\Lambda \setminus \hat{\Lambda}^\alpha}|\mathbf{Y}_{\Lambda \setminus \hat{\Lambda}^\alpha}) \\ & \geq \bar{L}(\tilde{\mathbf{W}}_{\hat{\Lambda}^\alpha}, \tilde{\mathbf{W}}_{\partial \hat{\Lambda}^\alpha}|\mathbf{Y}_{\hat{\Lambda}^\alpha}) + L(\tilde{\mathbf{W}}_{\Lambda \setminus \hat{\Lambda}^\alpha}|\mathbf{Y}_{\Lambda \setminus \hat{\Lambda}^\alpha}) + \bar{H}^\alpha(\tilde{\mathbf{W}}_{\hat{\Lambda}^\alpha}, \tilde{\mathbf{W}}_{\partial \hat{\Lambda}^\alpha}|\mathbf{Y}_{\hat{\Lambda}^\alpha}) + H^\alpha(\tilde{\mathbf{W}}_{\Lambda \setminus \hat{\Lambda}^\alpha}|\mathbf{Y}_{\Lambda \setminus \hat{\Lambda}^\alpha}), \end{aligned} \tag{2.13}$$

where

$$\begin{aligned} \bar{L}(\mathbf{W}_{\hat{\Lambda}^\alpha}, \mathbf{W}_{\partial \hat{\Lambda}^\alpha}|\mathbf{Y}_{\hat{\Lambda}^\alpha}) &= \beta \sum_{\langle ij \rangle, i, j \in \hat{\Lambda}^\alpha} \mathbf{1}_{\{w_i = w_j\}}(\mathbf{W}) + \beta \sum_{\langle ij \rangle, i \in \hat{\Lambda}^\alpha, j \in \partial \hat{\Lambda}^\alpha} \mathbf{1}_{\{w_i = x_j\}}(\mathbf{W}) \\ &+ h \sum_{i \in \hat{\Lambda}^\alpha} \mathbf{1}_{\{w_i = y_i\}}(\mathbf{W}), \end{aligned}$$

and $\partial \hat{\Lambda}^\alpha = \{j \in \Lambda \setminus \hat{\Lambda}^\alpha: \langle ij \rangle \text{ for some } i \in \hat{\Lambda}^\alpha\}$ (and analogously for \bar{H}^α).

Since outside $\hat{\Lambda}^\alpha$ the configurations $\tilde{\mathbf{W}}$ and \mathbf{W} coincide, the second and last terms of both sides of inequality (2.13) simplify. Furthermore, analogously to equation (2.12),

$$\bar{H}^\alpha(\mathbf{W}_{\hat{\Lambda}^\alpha}, \mathbf{W}_{\partial \hat{\Lambda}^\alpha}|\mathbf{Y}_{\hat{\Lambda}^\alpha}) = 0.$$

Hence inequality (2.13) reduces to

$$\bar{L}(\mathbf{W}_{\hat{\Lambda}^\alpha}, \mathbf{W}_{\partial \hat{\Lambda}^\alpha}|\mathbf{Y}_{\hat{\Lambda}^\alpha}) - \bar{L}(\tilde{\mathbf{W}}_{\hat{\Lambda}^\alpha}, \tilde{\mathbf{W}}_{\partial \hat{\Lambda}^\alpha}|\mathbf{Y}_{\hat{\Lambda}^\alpha}) \geq \bar{H}^\alpha(\tilde{\mathbf{W}}_{\hat{\Lambda}^\alpha}, \tilde{\mathbf{W}}_{\partial \hat{\Lambda}^\alpha}|\mathbf{Y}_{\hat{\Lambda}^\alpha}).$$

But the right-hand term is by definition non-negative, i.e.

$$\bar{L}(\mathbf{W}_{\hat{\Lambda}^\alpha}, \mathbf{W}_{\partial\hat{\Lambda}^\alpha} | \mathbf{Y}_{\hat{\Lambda}^\alpha}) \geq \bar{L}(\tilde{\mathbf{W}}_{\hat{\Lambda}^\alpha}, \tilde{\mathbf{W}}_{\partial\hat{\Lambda}^\alpha} | \mathbf{Y}_{\hat{\Lambda}^\alpha}).$$

However, since $\tilde{\mathbf{W}}_{\hat{\Lambda}^\alpha}$ has not been fixed and $\tilde{\mathbf{W}}_{\partial\hat{\Lambda}^\alpha} = \mathbf{W}_{\partial\hat{\Lambda}^\alpha}$, this is equivalent to

$$\bar{L}(\mathbf{W}_{\hat{\Lambda}^\alpha}, \mathbf{X}_{\partial\hat{\Lambda}^\alpha} | \mathbf{Y}_{\hat{\Lambda}^\alpha}) \geq \bar{L}(\mathbf{X}_{\hat{\Lambda}^\alpha}, \mathbf{X}_{\partial\hat{\Lambda}^\alpha} | \mathbf{Y}_{\hat{\Lambda}^\alpha}), \quad \forall \mathbf{X} \in \Omega, \quad \forall \mathbf{W} \in \mathcal{W}^\alpha, \quad (2.14)$$

i.e. in the set $\hat{\Lambda}^\alpha$, for any boundary conditions fixed in $\partial\hat{\Lambda}^\alpha$, the maximum of $\bar{L}(\mathbf{X} | \mathbf{Y})$ is realized by colouring the pixels in $\hat{\Lambda}^\alpha$ with the colour α .

Take any \mathbf{Z} in the set of configurations that realizes the global maximum (i.e. satisfies inequality (2.11)). There are two possibilities.

- (a) $\mathbf{Z} \in \mathcal{W}^\alpha$: in this case the theorem is proven.
- (b) For some $j \in \hat{\Lambda}^\alpha, \mathbf{Z}_j \neq \alpha$: in this case we define $\hat{\mathbf{W}} \in \mathcal{W}^\alpha$ sitewise as

$$\hat{W}_i = \begin{cases} \alpha & \text{for } i \in \hat{\Lambda}^\alpha, \\ \mathbf{Z}_i & \text{for } i \in \Lambda \setminus \hat{\Lambda}^\alpha. \end{cases}$$

Since $\hat{\mathbf{W}} = \mathbf{Z}$ on $\Lambda \setminus \hat{\Lambda}^\alpha$, by the definition of $\hat{\mathbf{W}}$,

$$\begin{aligned} L(\hat{\mathbf{W}} | \mathbf{Y}) &= \bar{L}(\hat{\mathbf{W}}_{\hat{\Lambda}^\alpha}, \hat{\mathbf{W}}_{\partial\hat{\Lambda}^\alpha} | \mathbf{Y}_{\hat{\Lambda}^\alpha}) + L(\hat{\mathbf{W}}_{\Lambda \setminus \hat{\Lambda}^\alpha} | \mathbf{Y}_{\Lambda \setminus \hat{\Lambda}^\alpha}) \\ &= \bar{L}(\hat{\mathbf{W}}_{\hat{\Lambda}^\alpha}, \mathbf{Z}_{\partial\hat{\Lambda}^\alpha} | \mathbf{Y}_{\hat{\Lambda}^\alpha}) + L(\mathbf{Z}_{\Lambda \setminus \hat{\Lambda}^\alpha} | \mathbf{Y}_{\Lambda \setminus \hat{\Lambda}^\alpha}) \\ &\geq \bar{L}(\mathbf{Z}_{\hat{\Lambda}^\alpha}, \mathbf{Z}_{\partial\hat{\Lambda}^\alpha} | \mathbf{Y}_{\hat{\Lambda}^\alpha}) + L(\mathbf{Z}_{\Lambda \setminus \hat{\Lambda}^\alpha} | \mathbf{Y}_{\Lambda \setminus \hat{\Lambda}^\alpha}) \end{aligned}$$

by inequality (2.14). But the last line of the above equation equals $L(\mathbf{Z} | \mathbf{Y})$; hence $\hat{\mathbf{W}} \in \mathcal{W}^\alpha$ is also a global maximum. □

The theorem is not too surprising. Intuitively, if the colour α is the correct colour in the coded problem (when the two colours are merged), or, in other words, if for a pixel the best possible choice is the colour α , compared with the other two undistinguished colours, then colour α will also be the best choice (up to equivalent MAP estimates) in the three-colour problem (when the two colours are not merged). Or, if colour 0, say, is preferred to colours 1 and 2 merged together and thus is stronger, then 0 will also be preferred when 1 and 2 are each on their own. Of course, owing to non-uniqueness of the posterior mode, there may be other MAP estimators with different site colourings, but with the same posterior probability.

If $Z_i^\alpha = 0$ then α is the true MAP restoration, for $\alpha \in \{0, 1, 2\}$. However, it is possible that $Z_i^\alpha = 1$ for all $\alpha \in \{0, 1, 2\}$. (It is easy to see that there is no further possible outcome. If $Z_i^0 = 0$, say, then necessarily $Z_i^1 = Z_i^2 = 1$ since in both cases the colour 0 will be merged with a second colour and will hence continue to be the optimal choice.) In this case the colour of the pixel will remain undecided and we shall denote it by ?.

To summarize our algorithm works as follows. We consider three codings

$$\begin{aligned} 0 &\rightarrow 0, & 1, 2 &\rightarrow 1, \\ 1 &\rightarrow 0, & 0, 2 &\rightarrow 1, \\ 2 &\rightarrow 0, & 0, 1 &\rightarrow 1 \end{aligned}$$

of the data. For each we run a Ford-Fulkerson algorithm to obtain $\mathbf{Z}^0, \mathbf{Z}^1$ and \mathbf{Z}^2 . Finally we define $\mathbf{X}^* \in \{0, 1, 2, ?\}^\Lambda$ as

$$X_i^* = \begin{cases} \alpha & \text{if } Z_i^\alpha = 0 \text{ for one } \alpha \in \{0, 1, 2\}, \\ ? & \text{otherwise.} \end{cases}$$

Could it happen that $Z_i^\alpha = 0$ for more than one α ? First, it is a consequence of the proof that any such colouring would yield the same posterior probability; hence the colour of X_i^* can be put equal to any of these α s. Furthermore, a quite tedious proof that we omit allows us to exclude this possibility completely.

Our theorem says that $X_i^* = \hat{X}_i$ whenever $X_i^* \neq ?$. The approximation of $\hat{\mathbf{X}}$ obtained by means of \mathbf{X}^* consists of the non-restored ?s in \mathbf{X}^* ; note that \mathbf{X}^* is computed in $O(C|\Lambda|^2)$ steps, i.e. in polynomial time in $|\Lambda|$.

3. EXPECTED NUMBER OF QUESTION-MARKS

The number of ?s left in \mathbf{X}^* measures the quality of the approximation. Our algorithm is deterministic and the output \mathbf{X}^* depends only on the input \mathbf{Y} . Hence the number of ?s in \mathbf{X}^* depends on the data \mathbf{Y} , and we define

$$N_\Lambda(\mathbf{Y}) = |\{i \in \Lambda: X_i^* = ?, \mathbf{X}^* = \mathbf{X}^*(\mathbf{Y})\}|.$$

In Section 5 we show several experiments with our method and we always compute the corresponding $N_\Lambda(\mathbf{Y})$ or the relative approximation $N_\Lambda(\mathbf{Y})/|\Lambda|$. In this section we estimate $N_\Lambda(\mathbf{Y})$ theoretically. From Section 2.1 we see that \mathbf{Y} is a stochastic process, obtained from equations (2.3) and (2.2). Hence $N_\Lambda(\mathbf{Y})$ is a random variable. We denote by $P_{\epsilon, \beta}(\mathbf{Y})$ the probability of \mathbf{Y} obtained from equations (2.2) and (2.3) and we want to estimate the expected value of $N_\Lambda(\mathbf{Y})$,

$$E_{\epsilon, \beta}\{N_\Lambda(\mathbf{Y})\} = \sum_{k=0}^{|\Lambda|} k P_{\epsilon, \beta}(\mathbf{x} \in \Omega: N_\Lambda(\mathbf{x}) = k).$$

Here is our result.

Theorem 2. There are two constants $c = c(\epsilon, \beta)$ and $C = C(\epsilon, \beta)$, strictly positive and independent of $|\Lambda|$, such that

$$c \leq \lim_{\Lambda \uparrow \mathbb{Z}^2} \left[\frac{E_{\epsilon, \beta}\{N_\Lambda(\mathbf{Y})\}}{|\Lambda|} \right] \leq C.$$

Proof of lower bound. Assume that $h < 4\beta$. (If $h \geq 4\beta$ then $X^* = Y$.) Consider a box B consisting of b^2 sites. Consider a configuration $\bar{\mathbf{Y}}_B$ in the $\{0, 1, 2\}^B$ -space such that the following property holds: for any $\bar{\mathbf{Y}}$ coinciding with $\bar{\mathbf{Y}}_B$ in B ,

there is a site i in B such that $X_i^* = ?$ when \mathbf{X}^* is obtained from $\bar{\mathbf{Y}}$.

Such a box always exists. For instance consider the case of nearest neighbour interactions and a square B centred on the origin with b^2 sites (b odd). Let

$$B_1 = \{(i, j): (i > 0 \text{ and } j \geq 0) \text{ or } (i \leq 0 \text{ and } j > 0)\} \cap B,$$

$$B_2 = \{(i, j): (i \geq 0 \text{ and } j < 0) \text{ or } (i < 0 \text{ and } j \leq 0)\} \cap B.$$

Consider the configuration $\bar{\mathbf{Y}}_B$,

$$Y_{i,j} = \begin{cases} 1 & \text{if } (i, j) \in B_1, \\ 2 & \text{if } (i, j) \in B_2, \\ 0 & \text{if } (i, j) = (0, 0). \end{cases}$$

Now choose b sufficiently large such that, *independently of the configuration \bar{Y} outside B* , $\hat{\Lambda}^1 = B_1$ and $\hat{\Lambda}^2 = B_2$, where $\hat{\Lambda}^\alpha$ is defined at the beginning of the proof of theorem 1. This is possible for any h and β because in the $\alpha = 1$ coded problem (1 against 0 and 2) the number of terms in equation (2.6) that multiply h is $|B_1|$ which grows roughly as $|\partial B_1|^2$, whereas the number of terms that multiply β is of the order $|\partial B_1|$. Hence, for any b sufficiently large, the best choice is to put 1s in B_1 . The same argument works for the $\alpha = 2$ coded problem (2 against 0 and 1). Furthermore, it is easy to see that if $h < 4\beta$ we obtain a ? at the origin. Hence the above configuration \bar{Y}_B has the following property:

$$X_0^* = ? \text{ when } \mathbf{X}^* \text{ is obtained from } \bar{Y}.$$

Let $p = \min_{Z_{B^c}} \{P_{\epsilon, \beta}(\bar{Y}_B | Z_{B^c})\}$, where the minimum is taken over all possible boundary conditions Z_{B^c} . Divide the lattice Λ into non-overlapping boxes of the same type as B . If \mathbf{Y} coincides with \bar{Y}_B on any such box, then our algorithm will produce one ? in the centre of the box. Altogether there are $|\Lambda|/b^2$ boxes.

Hence $N_\Lambda(\mathbf{Y})$ is bounded stochastically from below by a binomial random variable with parameters p and $|\Lambda|/b^2$, so that

$$E_{\epsilon, \beta}\{N_\Lambda(\mathbf{Y})\} \geq p|\Lambda|/b^2$$

and $c = p/b^2$.

Proof of upper bound. We have

$$E_{\epsilon, \beta}\{N_\Lambda(\mathbf{Y})\} = E_{\epsilon, \beta}\left(\sum_{i \in \Lambda} \mathbf{1}_{\{X_i^* = ? | \mathbf{Y}\}}\right).$$

If we can show that

$$P_{\epsilon, \beta}(\mathbf{Y} \in \Omega: X_i^* = ?) \leq C$$

for C independent of $|\Lambda|$, then we have the claim

$$E_{\epsilon, \beta}\{N_\Lambda(\mathbf{Y})\} \leq C|\Lambda|.$$

We need to find the upper bound.

As before in this proof we obtain a rough estimate of $P_{\epsilon, \beta}(\mathbf{Y} \in \Omega: X_i^* = ?)$. First consider again a box B with b^2 sites, where $b = b(\beta, \epsilon)$, and a configuration \bar{Y}_B such that

if \bar{Y}_B takes only one colour then \mathbf{X}_B^* is equal to that colour.

(It is easily seen that if $b > 4\beta/h$ then such a box is sufficiently large to have this property.) Next take the site i to represent the centre of the box B . We have

$$P_{\epsilon, \beta}(\mathbf{Y} \in \Omega: X_i^* \neq ?) \geq P_{\epsilon, \beta}(\bar{Y}_B) =: q$$

which is independent of $|\Lambda|$. It follows that $C = 1 - q$. □

From the proof it can be seen that good bounds can be obtained by specific choices of B and \bar{Y}_B . In the lower bound the choice should make p/b^2 as large as possible. In the upper bound the proof can be improved by considering more ?-free events of the type \bar{Y}_B . Note that we have decided to use a very special configuration where we can check locally that a ? will be produced. Of course

this happens in general also because of global effects.

Rough estimates are easy to obtain. From a practical point of view they are not useful, however. To obtain better estimates for p and q , we should consider as many configurations \overline{Y}_B as possible. To obtain an idea of their optimal values we consider the one-dimensional case for given β , ϵ and b . We can then enumerate all configurations giving rise to a $?$. Let us take $b = 5$ and β and ϵ such that $0.5 < \beta/h < 1$. Writing out the $3^5 = 243$ possible configurations on the box B , we find 54 configurations which give rise to a $?$ at the centre of the box B , independently of the boundary conditions on the box. This implies that in this case $C = 54/243 \sim 0.2$.

4. OTHER MODELS

We next generalize our algorithm to images generated by the Potts model with a greater number of colours. This generalization is not straightforward, however, as it involves the introduction of a further step in the algorithm, which we call 'cascade coding', and which we describe briefly for four-colour images for simplicity.

The first step, when applied to four-colour images, is simply the extension of the coding transformation (2.7) by considering $\alpha \in \{0, 1, 2, 3\}$. This produces a first approximation $\mathbf{X}^* \in \{0, 1, 2, 3, ?\}^A$ to the MAP estimate. We are assured (by a simple extension of theorem 1) that the colours attributed are exactly those of the MAP estimate. We are left with the $?$ -clusters, but this time we can perform a second coding, of the pixels belonging to these $?$ -clusters only, to try to solve some of them. The coding is

$$\begin{aligned} 0, 1 \rightarrow 0, & \quad 2, 3 \rightarrow 1, \\ 0, 2 \rightarrow 0, & \quad 1, 3 \rightarrow 1, \\ 0, 3 \rightarrow 0, & \quad 1, 2 \rightarrow 1. \end{aligned}$$

As previously, if we apply the Ford-Fulkerson algorithm to these coded images, we shall obtain $\mathbf{Z}^k \in \{0, 1\}^A$, $k \in \{1, 2, 3\}$ as solutions to the three two-colour maximization problem. We must next define a decoding algorithm to obtain, from \mathbf{Z}^k , information on the four-colour maximization problem over the $?$ -clusters.

We conjecture that the decoding algorithm should run as follows. If for a site i we have $Z_i^k = 0$, $\forall k \in \{1, 2, 3\}$, then $\hat{X}_i = 0$. If, however, we have $Z_i^1 = 0$ and $Z_i^2 = Z_i^3 = 1$, then $\hat{X}_i = 1$, and similarly for $Z_i^2 = 0$, $Z_i^1 = Z_i^3 = 1$ and $Z_i^3 = 0$, $Z_i^1 = Z_i^2 = 1$, which would give respectively $\hat{X}_i = 2$ and $\hat{X}_i = 3$. All other outcomes would give rise to 'double question-marks' (??). The reason is the same as before: if a colour is always preferred, helped by any of the other three, it is also preferred in the full four-colour MAP estimate.

In other words, in the four-colour case we must not only test the 'strength' of any given colour against the three other colours merged together but also how colours behave when merged two by two. If we venture even further, it becomes apparent that for five colours we have essentially the same situation (one colour against the other four, then two colours against the other three) but that for six colours, for example, we must consider yet another step (colours merged three

by three). The algorithm becomes increasingly more complicated but remains polynomial in $|\Lambda|$. Also only increasingly smaller clusters must be treated.

Our algorithm works, for the Potts model, not only with nearest neighbour interactions but also when we consider the contribution of the sites 'on the diagonals', at distance $\sqrt{2}$. It is easy to extend theorem 1 to that case, as the number of neighbours taken into account when calculating the posterior energy (2.6) is of no consequence in the proof presented here. More generally, only pairwise interactions are required.

In many applications the noise follows a continuous distribution, say pixelwise independent Gaussian. We cannot operate a coding of such data. Of course, if the noise can be discretized to produce the same number of colours as in the original image, our method would be practicable.

5. SIMULATION STUDY

We present some of the results of the computer simulations that we have carried out on 64×64 pixels. To study the typical number of ?s produced by the coding and decoding algorithm for various values of the noise parameter ϵ , we first produced a 'clean' image with a Gibbs sampler, an example of which is shown in Fig. 1(a). We then corrupted the clean image with noise (see for example Fig. 1(b) with $\epsilon = 0.2$) and applied the coding and decoding algorithm to this 'dirty' image, to produce an image containing a new, fourth, colour—representing the ?—such as shown in Fig. 1(c). Finally, for comparison, we also applied the simulated annealing algorithm to the corrupted image (see for example Fig. 1(d)), to be able to judge the efficiency of our algorithm against other methods used so far.

We show four groups of images, all of which were generated starting from Fig. 1(a) with increasing values of the noise parameter ϵ . From Figs 1(c), 2(b), 3(b) and the results given in Table 1 we can see that, in general, the number of ?s produced by the algorithm increases as the value of the noise parameter increases. And although, for a sufficiently high level of noise, the approximation to the MAP estimate produced by the algorithm may consist *entirely* of ?s (see Fig. 4(b)), if the value of ϵ is not too large, the number of ?s produced is limited. Moreover, it is easy to see that the ?s tend to be localized on the interface between two regions having a definite colour, or in regions which correspond, in the corrupted image, to areas where all three possible colours are present together. In other words, the ?s produced by the coding and decoding algorithm indicate regions of the image in which the data are the least reliable. Furthermore the ?-clusters are often small, making further estimation of the corresponding pixels easy.

We computed the percentage of ?s in the images produced by our method for various values of β and ϵ . These are displayed in the third column of Table 1. Note that the ?s increase when ϵ increases and decrease when β increases. Values of β over 1.2 make the picture smoother and are thus more relevant. Here the number of ?s is very small. Of course, as already mentioned, by increasing the noise we can produce a restoration displaying *only* ?s.

The next column in Table 1 is related to the performance of the simulated annealing algorithm run with a geometric schedule (i.e. the temperature decreases like 0.9^k where k is a full updating step: we stopped at temperature 0.001 starting

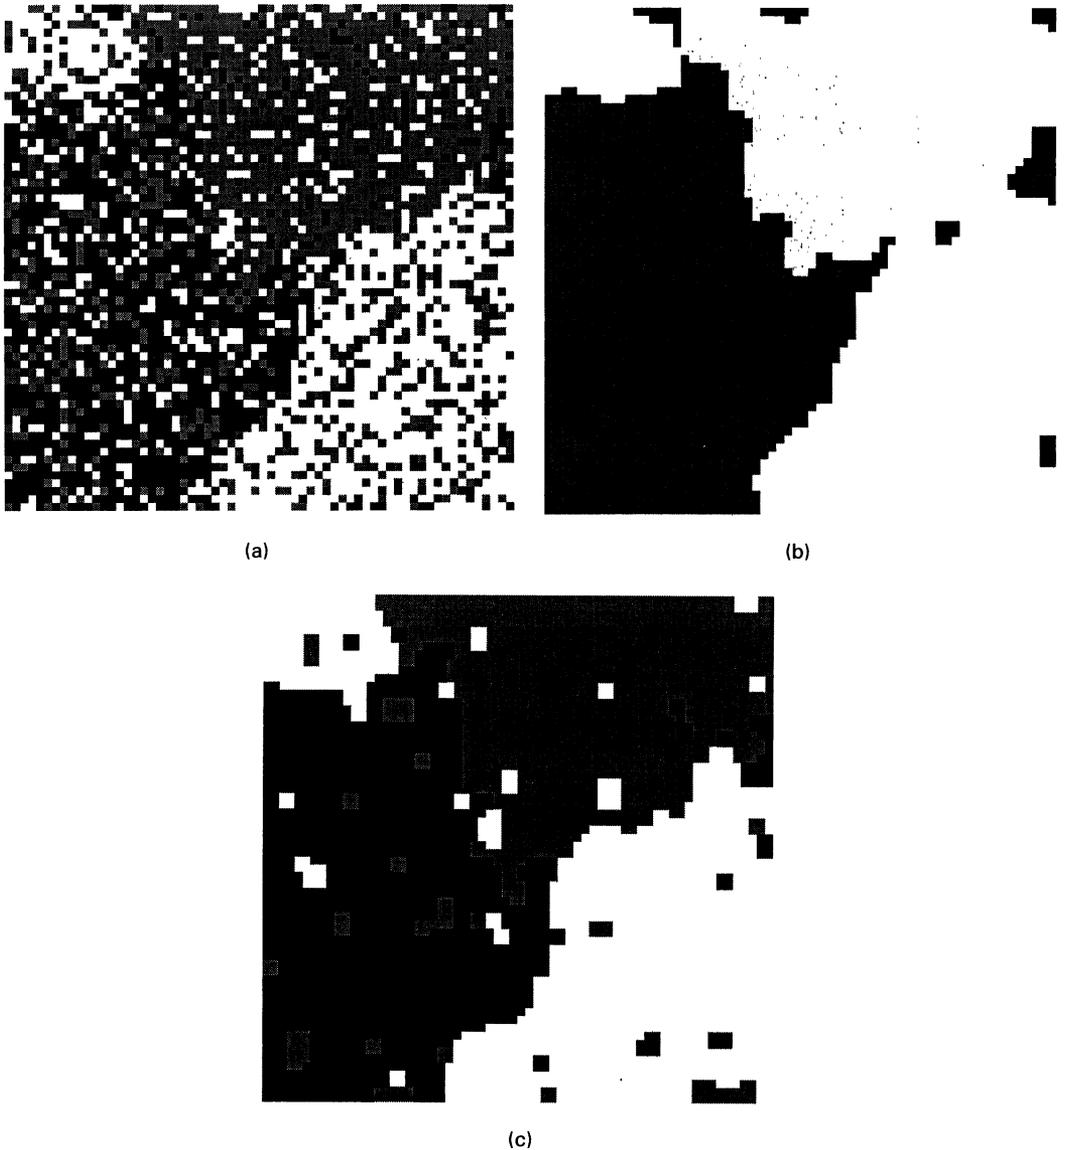


Fig. 2. (a) Image in Fig. 1(a) degraded with $\epsilon = 0.3$; (b) approximate MAP restoration of (a) (?s are shown in black, whereas other colours take grey levels); (c) simulated annealing (with geometric schedule) restoration of (a)

from 1). The percentage displayed is that of misclassified pixels with respect to the exact MAP estimate produced by our algorithm, excluding of course ?s. In the areas where our method produces the exact MAP estimate, simulated annealing makes mistakes, sometimes of the same order as the number of ?s. Starting from our three colour plus ? restorations we can compute the exact MAP estimate by full enumeration of the ?-clusters. If these are too big we could run simulated annealing

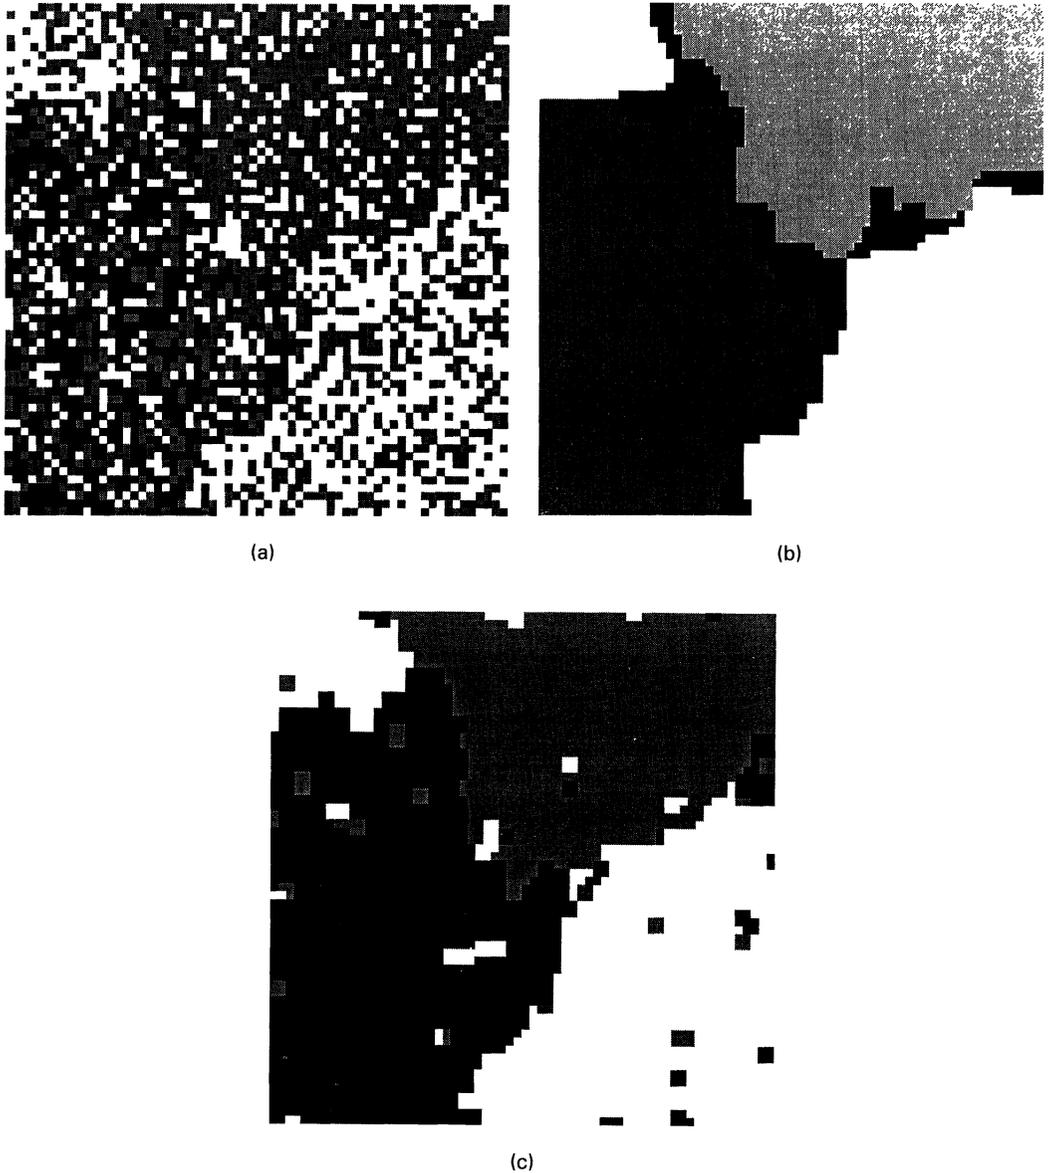


Fig. 3. (a) Image in Fig. 1(a) degraded with $\epsilon = 0.4$; (b) approximate MAP restoration of (a) (?'s are shown in black, whereas other colours take grey levels); (c) simulated annealing (with geometric schedule) restoration of (a)

on them (with the correct boundary conditions) with a logarithmic schedule. With respect to this exact MAP estimate, simulated annealing on the whole scene (with geometric schedule) still performs reasonably. If we compare this final column with the third, i.e. the percentage of ?'s and percentage of misclassified pixels in simulated annealing, we observe that for high values of β our method is better than or equivalent to simulated annealing. For low values of β our method produces more

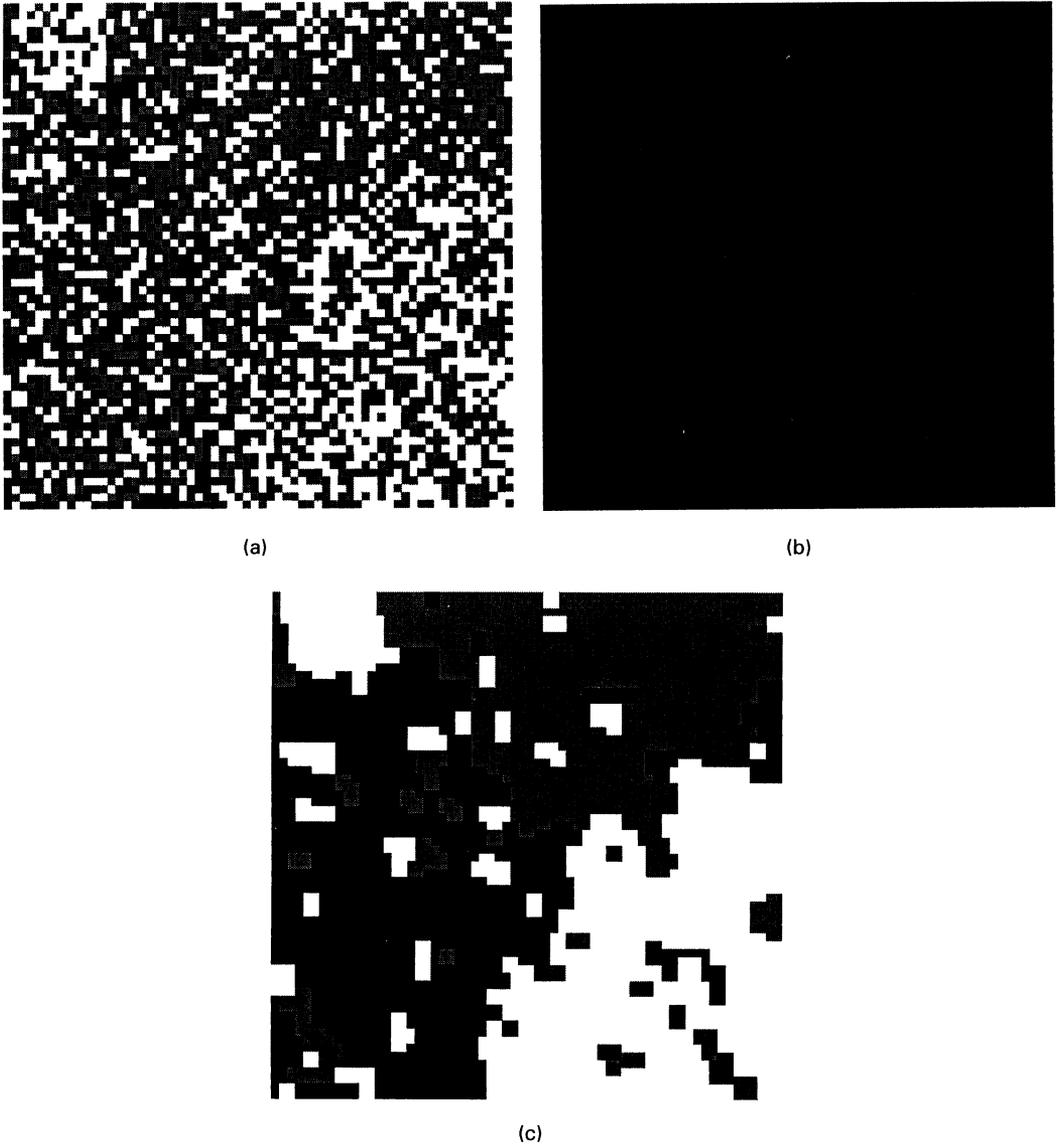


Fig. 4. (a) Image in Fig. 1(a) degraded with $\epsilon = 0.5$; (b) approximate MAP restoration of (a) (all pixels are ?s and are shown in grey); (c) simulated annealing (with geometric schedule) restoration of (a)

?s. But remember the two completely different computational burdens. Moreover the time required to produce an exact MAP restoration (as described here) is less than that for simulated annealing with a logarithmic schedule.

ACKNOWLEDGEMENTS

This research was started when AF was visiting the University of São Paulo, whose hospitality is here acknowledged. PGS was supported by an Erasmus grant

TABLE 1

| β | ϵ | <i>New algorithm:</i> % of ?s | <i>Simulated annealing:</i> % of misclassification with respect to exact MAP restoration excluding ?-areas | <i>Simulated annealing:</i> % of misclassification with respect to full exact MAP restoration |
|---------|------------|----------------------------------|---|--|
| 1.4 | 0.2 | 0.1 | 1.1 | 1.1 |
| | 0.3 | 0.0 | 3.4 | 3.4 |
| | 0.4 | 0.0 | 7.3 | 7.3 |
| 1.2 | 0.2 | 3.1 | 0.9 | 1.6 |
| | 0.3 | 4.4 | 3.6 | 3.9 |
| | 0.4 | 7.6 | 7.0 | 7.6 |
| | 0.5 | 100.0 | — | — |
| 1.0 | 0.1 | 6.9 | 0.9 | 1.2 |
| | 0.2 | 17.5 | 2.6 | 5.7 |
| | 0.3 | 35.6 | 3.1 | 6.1 |
| | 0.4 | 72.9 | 3.1 | 14.1 |
| 0.7 | 0.3 | 42.0 | 3.0 | 10.3 |

of the European Economic Community during her stay at the Istituto per le Applicazioni del Calcolo. Financial support was also provided by the Fundação Amparo Pesquisa Estado São Paulo and the Conselho Nacional Desenvolvimento Científico Tecnológico.

REFERENCES

- Aarts, E. and Korst, J. (1989) *Simulated Annealing and Boltzmann Machines*. New York: Wiley.
- Barahona, F. (1982) On the computational complexity of Ising spin-glass models. *J. Phys. A*, **15**, 3241–3248.
- Garey, M. R. and Johnson, D. S. (1979) *Computers and Intractability: a Guide to the Theory of NP-completeness*. San Francisco: Freeman.
- Geman, D. (1990) Random fields and inverse problems in imaging. *Lect. Notes Math.*, **1470**.
- Gidas, B. (1991) Metropolis type Monte Carlo simulation algorithms and simulated annealing. *Trends Contemp. Probab.*, to be published.
- Greig, D. M., Porteous, B. T. and Seheult, A. H. (1989) Exact maximum *a posteriori* estimation for binary images. *J. R. Statist. Soc. B*, **51**, 271–279.
- Jubb, M. D. (1989) Image reconstruction. *PhD Thesis*. University of Bath, Bath.