

A STUDY OF INTERACTIVE MULTIPLE CLASS IMAGE SEGMENTATION PROBLEMS

MICHAEL K. NG*, GUOPING QIU[†], AND ANDY M. YIP[‡]

Abstract. In this paper, we study a bilaterally constrained optimization model for the semi-supervised multiple class image segmentation problem. We prove that the solution of the corresponding unconstrained problem satisfies a discrete maximum principle. This implies that the bilateral constraints are satisfied automatically and that the solution is unique. We study domain decomposition methods for solving linear systems arising from the optimality conditions of the problem. We demonstrate that domain decomposition with block relaxation methods are quite effective and outperform other tested methods. We also carry out a numerical study on the effect of boundary conditions on the optimization problems which provides some insights into the specification of boundary conditions as an input knowledge in the learning context.

Key words. Image segmentation, image matting, discrete maximum principle, domain decomposition, boundary conditions

1. Introduction. Foreground-background segmentation has wide applications in computer vision (e.g. scene analysis), computer graphics (e.g. image editing) and medical imaging (e.g. organ segmentation). Fully automatic image segmentation has many intrinsic difficulties and is still a very hard problem. In many applications, such as image editing in computer graphics and organ segmentation in medical imaging, semi-automatic and interactive approaches, where human operators provide strong priors for the computational algorithms to perform segmentation, cannot only overcome the inherent technical difficulties of fully automatic image segmentation, but may also be desirable because the segmentation process and results can be controlled in many of these applications.

There have been increasing activities in the research community to develop interactive semi-automatic image segmentation techniques [1, 2, 3, 4, 5, 6]. In [1], Boykov and Jolly presented an interactive image segmentation technique based on graph cut. Users labeled seed pixels which indicating certain background and foreground were used as strong priors for segmenting images into figure and ground. In [2], Boykov and Kolmogorov showed that graph cut based segmentation algorithms could be implemented very fast. In [5], Yu and Shi presented a segmentation given partial grouping constraints method. User inputs were used as bias to a natural grouping process, and they formulated such biased grouping problem as a constrained optimization problem that propagates sparse partial grouping information to the unlabeled data by enforcing grouping smoothness and fairness on the labeled data points. They used the normalized cut criterion and solved the optimization problem by eigen-decomposition. In [3], Rother *et al.* presented an interactive image foreground extraction method that was computationally based on graph cut but they introduced a simpler user interaction technique to reduce user efforts in the interaction process and an iterative model updating procedure to improve accuracy. In [4], an interactive foreground background segmentation method was introduced in the context of image matting. Wang and

*Centre for Mathematical Imaging and Vision and Department of Mathematics, Hong Kong Baptist University, Kowloon Tong, Hong Kong (mng@math.hkbu.edu.hk). Research supported in part by Hong Kong Research Grants Council Grant Nos. 7035/04P and 7035/05P, and HKBU FRGs.

[†]School of Computer Science, University of Nottingham, United Kingdom (qiu@cs.nott.ac.uk).

[‡]Department of Mathematics, National University of Singapore, 2, Science Drive 2, S117543, Singapore (andyyip@nus.edu.sg).

Cohen used belief propagation to iteratively propagate user labeled pixels to the unlabeled pixels. Recently, Guan and Qiu [6] have developed an optimization based figure ground segmentation technique, where a transparency image was computed by optimizing a quadratic cost function with user supplied linear constraints.

1.1. Image Matting Problems. The observed image $u = (u_i)$ is modeled as a convex combination of a foreground $F = (F_i)$ and a background $B = (B_i)$:

$$u_i = \alpha_i F_i + (1 - \alpha_i) B_i.$$

Here i denotes a linear order of the pixel location and $\alpha = (\alpha_i)$ with $0 \leq \alpha_i \leq 1$ for each i indicates the degree of membership of each pixel to the foreground. When $\alpha_i = 1$, the i -th pixel is a certain foreground pixel. When $\alpha_i = 0$, the pixel is a certain background pixel. The problem of image matting is to estimate F , B and α from the given image u . Once F , B and α are obtained, we can replace the original background B by another totally different image to synthesize a new image. The result is a new image having the same object but with a different background. While a hard segmentation i.e. $\alpha_i \in \{0, 1\}$ for each i can be used, the resulting synthesized image often looks unnaturally especially when the object contains hair. The idea of image matting is to allow α to take a continuum between 0 and 1, presumably at a thin layer around the interface between the foreground and background. In this way, the pixels close to the interface can be synthesized as a blending of the foreground and background, thus looking more natural.

The image matting problem is seriously under constrained where we have one equation and three unknown variables. In [6], Guan and Qiu regarded the alpha value in each pixel as the probability of the pixel belonging to the foreground and formulated the problem of computing the alpha matte as a constrained optimization problem. If two pixels both belong to the foreground (or background), it is reasonable to assume that they will have similar colors and other photometric properties. In other words, pixels with similar photometric properties should have similar α values. Within the background or foreground, it is also reasonable to assume that the image is “smooth”, i.e., neighboring pixels will have similar photometric properties hence α values. If two pixels one belongs to the background, the other belongs to the foreground, it is also reasonable to assume that they will have very different photometric properties, that is, two pixels with different photometric properties should have different α values. Based on such reasoning, Guan and Qiu formulated a cost function for the transparency image $\alpha(z)$ such that an $\alpha(z)$ that satisfies the above assumptions will correspond to a minimum of the cost function. Note that they only computed the alpha matte and did not estimate the foreground and background values in the matting equation but instead they performed figure and ground segmentation by thresholding the alpha matte. In this paper, we will follow the same practice and will use matting and segmentation interchangeably, both should be understood as computing the transparency image α and thresholding α achieves crispy segmentation.

For multiple class image segmentation [7], multiple α 's can be introduced each corresponds to the probability of the pixels belonging to one class of object/region [6]. A cost function that is similar to that of the case of figure and ground segmentation can also be formulated. Unlike previous approaches, e.g., [10, 1], where multiple class segmentation had to be done by recursively applying two class (figure and ground) segmentation, this approach formulated multiple class segmentation in a single cost function and a solution to this cost function achieves multiple class segmentation in a

single step. However, optimizing the multiple class cost function is computationally much more challenging [7].

1.2. Outline. The main aim of this paper is to study a bilaterally constrained optimization model for the semi-supervised multiple class image segmentation problem. We prove that the solution of the corresponding unconstrained problem satisfies a discrete maximum principle. This implies that the bilateral constraints are satisfied automatically and that the solution is unique. We study domain decomposition methods for solving linear systems arising from the optimality conditions of the problem. We demonstrate that domain decomposition with block relaxation methods are quite effective and outperform other tested methods.

An important aspect of this and other similar interactive approaches, e.g., [1, 8], is that users have to supply constraints (manually labeled pixels). These user supplied constraints greatly affect the segmentation results. However, in the literature, there is no theory to guide users to label the most suitable pixels. We carry out a numerical study on the effect of boundary conditions on the optimization problems which provides some insights into the specification of boundary conditions as an input knowledge in the learning context; and such insights can provide guidance for the users to label suitable pixels.

The outline of this paper is given as follows. In Section 2, we study multiple class image matting problems. In Section 3, we analyze the linear system arising from the problem. In Section 4, we study numerical solvers for this linear system. In Section 5, numerical results are presented to illustrate the performance of numerical solvers. We also study the effect of boundary conditions in image matting problems. Finally, concluding remarks are given in Section 6.

2. Interactive Image Matting Problems.

2.1. Matting Neighbors. In a matting model, alpha mattes are based on the similarity of both geometric and photometric neighbors. They define the geometric neighborhood of pixel i as:

$$N_i^g := \{j \in \mathbb{Z}_+^2 : 0 < \|i - j\| \leq r_g\}$$

where $r_g > 0$ is a constant controlling the size of the neighborhood and $\|\cdot\|$ is a vector norm. For example, we often use a 3×3 window around a pixel as its geometric neighborhood (excluding pixel i itself). This corresponds to $r_g = 1$ and $\|\cdot\| = \|\cdot\|_\infty$.

Let G_i be a feature vector (e.g. color, texture, etc.) computed around pixel i . We call G_i a photometric feature. Then, a natural way to define photometric neighborhood N_i^p of the pixel i is pixels whose photometric features are close to G_i . But this may lead to a large neighborhood. Moreover, these neighbors are time-consuming to compute. Instead, we can define the photometric neighborhood N_i^p of i to be pixels that are the top k photometrically closest (in Euclidean norm) pixels to i within a window around i (excluding i itself). In our experiments, each feature vector has 27 entries which are the RGB values of the 3×3 window around i , the window size is set to 17×17 and the top 4 neighbors are used, see for instance [6].

The matting neighbor N_i is defined to be $N_i^p \cup N_i^g$ which will be used to define an alpha matte. The geometric neighbor is a symmetric relation, i.e. $j \in N_i^g$ if and only if $i \in N_j^g$. However, photometric neighbor, and hence matting neighbor, is non-symmetric since it is defined based on k nearest neighbors.

2.2. An Optimization Model for Two-class Problems. The basic idea of the optimization model is that the alpha matte of neighboring pixels (either geometric or photometric) should be similar. More precisely, the optimization model for the two-class matting problem is given by:

$$\min_{\alpha} \sum_{i \in \Omega} \left[\alpha_i - \sum_{j \in N_i} (w_{ij}^g + \lambda w_{ij}^p) \alpha_j \right]^2$$

subject to

$$0 \leq \alpha_i \leq 1, \quad \text{for all } i$$

and the boundary conditions

$$\alpha_i = \begin{cases} 1, & \text{on } \mathcal{F}, \\ 0, & \text{on } \mathcal{B}. \end{cases} \quad (2.1)$$

Here, $w_{ij}^g \geq 0$ is the geometric similarity score between pixel i and pixel j , $w_{ij}^p \geq 0$ is the photometric similarity score between pixel i and pixel j , $\lambda \geq 0$ is a user specified parameter controlling the relative importance of the geometric and photometric similarity scores. The sets \mathcal{F} and \mathcal{B} are pixels that are certain foreground and certain background respectively, supplied by the user interactively. These boundary conditions ensure the uniqueness of solution to the problem. We will prove such property in the next section. The set Ω consists of pixels at which the alpha matte is unknown, and its complement Γ is equal to $\mathcal{F} \cup \mathcal{B}$.

In the matting model, the similarity scores w_{ij}^g and w_{ij}^p are defined by:

$$w_{ij}^g := \begin{cases} c_i^g e^{-\|i-j\|_2^2 / \sigma_g^2}, & \text{if } j \in N_i^g, \\ 0, & \text{otherwise,} \end{cases}$$

$$w_{ij}^p := \begin{cases} c_i^p e^{-\|G_i - G_j\|_2^2 / \sigma_p^2}, & \text{if } j \in N_i^p, \\ 0, & \text{otherwise,} \end{cases}$$

where σ_g^2 is the variance of the geometric locations within N_i^g , σ_p^2 is the variance of the photometric features within N_i^p , c_i^g and c_i^p are normalizing constants so that $\sum_j w_{ij}^g = \sum_j w_{ij}^p = \frac{1}{1+\lambda}$. It can be easily seen that w_{ij}^g (respectively w_{ij}^p) measures the similarity between the geometric (respectively photometric) features of pixel i and pixel j .

Let $w_{ij} := w_{ij}^g + \lambda w_{ij}^p$. Then we have

$$\sum_j w_{ij} \equiv 1 \quad \forall i. \quad (2.2)$$

We note that w_{ij} is the relative combined geometric and photometric similarity measures of pixel i to other neighboring pixels.

Using the lexicographical ordering of the pixels, we can express the above optimization in matrix-vector notations:

$$\min_{\mathbf{0} \leq \alpha \leq \mathbf{1}} \|D_{\Omega}(I - W)\alpha\|_2^2 \quad (2.3)$$

subject to the boundary conditions (2.1). Here, $\mathbf{0}$ and $\mathbf{1}$ are vectors with all entries 0 and all entries 1 respectively, I is the identity matrix, $W = (w_{ij})$ is an $N \times N$ matrix

recording the weights, $\alpha = (\alpha_i)$ is an N vector representing the alpha matte, N is the number of pixels and D_Ω is an $|\Omega| \times N$ downsampling matrix from the image domain to Ω . The inequality constraints are understood entry-wise. The matrix W is non-symmetric because each pair of pixels have different sets of geometric and photometric neighbors, therefore, the normalization constants c_i^g and c_i^p are different for different pixels.

In [6], Guan and Qiu have presented experimental results on this model and showed that it is very effective in pulling the alpha matte from the image. Moreover, they showed that the use of both geometric and photometric neighbors is very important for its success. Furthermore, the image matting problem is often ambiguous. Therefore, they resolved this problem by requesting the user to supply pixels that are certain background and certain foreground.

2.3. An Optimization Model for Multiple-class Problems. In this subsection, we generalize the two-class model to the case of multiple-class image matting. This multi-class model allows us to handle images with multiple segments. Now the image is modeled as a convex combination of M images:

$$u_i = \alpha_i^1 F_i^1 + \alpha_i^2 F_i^2 + \dots + \alpha_i^M F_i^M.$$

Here, the M alpha mattes satisfy

$$0 \leq \alpha_i^m \leq 1$$

for each m and each i and

$$\sum_{m=1}^M \alpha_i^m = 1$$

for each i . In the multiple-class setting, the user supplies a set of M regions

$$\Gamma := \Gamma^1 \cup \Gamma^2 \cup \dots \cup \Gamma^M$$

where Γ^m is the set of pixels with certain membership to the m -th class. We assume that Γ is non-empty but each individual region Γ^m can be empty. We further assume that each pixel in Γ is a matting neighbor of some pixel in Ω . If the user specifies a large region which contains pixels that are not matting neighbor of Ω , then such pixels have no effect on the optimization problem. Therefore, we can safely ignore them.

The optimization model for pulling the alpha mattes is given as follows:

$$\min_{\alpha^1, \dots, \alpha^M} \sum_{m=1}^M \|D_\Omega(I - W)\alpha^m\|_2^2 \quad (2.4)$$

subject to $\mathbf{0} \leq \alpha^m \leq \mathbf{1}$, $\sum_m \alpha^m = \mathbf{1}$ and the boundary conditions

$$\alpha_i^m = \begin{cases} 1, & \text{on } \Gamma^m, \\ 0, & \text{on } \Gamma \setminus \Gamma^m, \end{cases} \quad (2.5)$$

for $m = 1, 2, \dots, M$. Similar to the two-class case, Ω is the set of pixels at which the alpha mattes are unknown (i.e. the complement of Γ) and D_Ω is the downsampling matrix from the image domain to Ω . When $M = 2$, this model can be shown to reduce to the two-class model in (2.3).

We observe that the objective and the constraints are convex. Moreover, the feasible set is compact. Therefore, a global minimizer exists. We will show in the next subsection that the solution is unique, see (2.9) and Corollary 3.3.

2.4. Optimality Conditions. For the sake of completeness, we derive in this subsection the optimality conditions for the multiple-class model (2.4). Indeed, we show that the optimization can be simplified so that the optimal alpha mattes α^m are computed independently. In contrast, Guan and Qiu [6] solved all the alpha mattes simultaneously which requires a huge computational cost, see the remarks at the end of this subsection. We here ignore the bilateral constraints $\mathbf{0} \leq \alpha^m \leq \mathbf{1}$. We will show in the next section that an explicit imposition of such constraints is indeed unnecessary, see Corollary 3.2.

Let us define some notations. Let $A = I - W$. Without loss of generality, we assume that the pixels are ordered so that A and α^m can be partitioned as follows:

$$A = \begin{bmatrix} A_{\Omega\Omega} & A_{\Omega\Gamma} \\ A_{\Gamma\Omega} & A_{\Gamma\Gamma} \end{bmatrix} \quad \text{and} \quad \alpha^m = \begin{bmatrix} \alpha_{\Omega}^m \\ \alpha_{\Gamma}^m \end{bmatrix}.$$

In fact, we have $A_{\Omega\Omega} = D_{\Omega}AD_{\Omega}^T$, $A_{\Omega\Gamma} = D_{\Omega}AD_{\Gamma}^T$ and $\alpha_{\Omega}^m = D_{\Omega}\alpha^m$ for $m = 1, 2, \dots, M$. Other variables with subscripts Ω or Γ are defined similarly.

The Lagrangian for (2.4) is given by

$$L(\alpha, \mu, \eta) = \frac{1}{2} \sum_{m=1}^M \|D_{\Omega}A\alpha^m\|_2^2 - \left\langle \sum_{m=1}^M \alpha_{\Omega}^m - \mathbf{1}, \mu_{\Omega} \right\rangle - \sum_{m=1}^M \langle \alpha_{\Gamma}^m - \bar{\alpha}_{\Gamma}^m, \eta_{\Gamma}^m \rangle$$

where $\bar{\alpha}_{\Gamma}^m$ are the given boundary data on Γ , $\mu_{\Omega} = (\mu_i)$ are the Lagrange multipliers for the constraints $\sum_m \alpha_{\Omega}^m = \mathbf{1}$ and $\eta_{\Gamma}^m = (\eta_i^m)$ are the Lagrange multipliers for the constraints $\alpha_{\Gamma}^m = \bar{\alpha}_{\Gamma}^m$. We have included the constant 1/2 in the first term for convenience.

The optimality conditions are given by

$$\frac{\partial L}{\partial \alpha^m} = \mathbf{0} \quad (m = 1, 2, \dots, M) \quad \text{and} \quad \frac{\partial L}{\partial \mu} = \frac{\partial L}{\partial \eta} = \mathbf{0}.$$

We will examine only the equation $\frac{\partial L}{\partial \alpha^m} = \mathbf{0}$ for the others just correspond to the equality constraints. We have

$$\frac{\partial L}{\partial \alpha^m} = A^T D_{\Omega}^T D_{\Omega} A \alpha^m - \begin{bmatrix} \mu_{\Omega} \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathbf{0} \\ \eta_{\Gamma}^m \end{bmatrix} = \mathbf{0}, \quad m = 1, 2, \dots, M. \quad (2.6)$$

By summing the above equation for all m and using the constraints $\sum_m \alpha^m = \mathbf{1}$, we have

$$A^T D_{\Omega}^T D_{\Omega} A \mathbf{1} = \begin{bmatrix} M\mu_{\Omega} \\ \sum_{m=1}^M \eta_{\Gamma}^m \end{bmatrix}, \quad m = 1, 2, \dots, M.$$

Notice that $A\mathbf{1} = (I - W)\mathbf{1} = \mathbf{1} - W\mathbf{1} = \mathbf{0}$ because of (2.2). Hence, we have $\mu_{\Omega} = \mathbf{0}$ and $\sum_{m=1}^M \eta_{\Gamma}^m = \mathbf{0}$. Equation (2.6) can now be written as:

$$A^T D_{\Omega}^T D_{\Omega} A \alpha^m = \begin{bmatrix} \mathbf{0} \\ \eta_{\Gamma}^m \end{bmatrix}, \quad m = 1, 2, \dots, M.$$

Using the partition of A and α^m , we have

$$\begin{bmatrix} A_{\Omega\Omega}^T A_{\Omega\Omega} & A_{\Omega\Omega}^T A_{\Omega\Gamma} \\ A_{\Gamma\Omega}^T A_{\Omega\Omega} & A_{\Gamma\Omega}^T A_{\Omega\Gamma} \end{bmatrix} \begin{bmatrix} \alpha_{\Omega}^m \\ \alpha_{\Gamma}^m \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \eta_{\Gamma}^m \end{bmatrix}, \quad m = 1, 2, \dots, M.$$

We will show in Corollary 3.4 that $A_{\Omega\Omega}$ is non-singular. We also notice that the boundary conditions are $\alpha_\Gamma^m = \bar{\alpha}_\Gamma^m$. Thus, the optimality conditions can be reduced to

$$A_{\Omega\Omega}\alpha_\Omega^m + A_{\Omega\Gamma}\alpha_\Gamma^m = \mathbf{0}, \quad (2.7)$$

$$\alpha_\Gamma^m = \bar{\alpha}_\Gamma^m. \quad (2.8)$$

Finally, we write the conditions as

$$\tilde{A}\alpha^m = \mathbf{b}^m, \quad m = 1, 2, \dots, M. \quad (2.9)$$

where

$$\tilde{A} = \begin{bmatrix} A_{\Omega\Omega} & A_{\Omega\Gamma} \\ 0 & I \end{bmatrix} \quad (2.10)$$

and

$$\mathbf{b}^m = \begin{bmatrix} \mathbf{0} \\ \bar{\alpha}_\Gamma^m \end{bmatrix}.$$

Here there are several remarks for the above equations.

1. \tilde{A} is defined by replacing a_{ij} with the Kronecker Delta δ_{ij} if the pixel i is in Γ .
2. The Lagrange multipliers η_Γ^m need not to be computed in order to obtain the optimal alpha mattes. But once the optimal alpha mattes are obtained, we can obtain η_Γ^m by $A_{\Omega\Gamma}^T A_{\Omega\Omega}\alpha_\Omega^m + A_{\Omega\Gamma}^T A_{\Omega\Gamma}\alpha_\Gamma^m$.
3. The optimal alpha mattes can be computed independently for each m . Therefore, reducing the computational complexity compared to the approach of Guan and Qiu which solves the least squares problem:

$$\min_{\alpha^1, \dots, \alpha^M} \left\| \begin{bmatrix} \tilde{A} & & & \\ & \tilde{A} & & \\ & & \ddots & \\ & & & \tilde{A} \\ I & I & \dots & I \end{bmatrix} \begin{bmatrix} \alpha^1 \\ \alpha^2 \\ \vdots \\ \alpha^M \end{bmatrix} - \begin{bmatrix} \mathbf{b}^1 \\ \mathbf{b}^2 \\ \vdots \\ \mathbf{b}^M \\ \mathbf{1} \end{bmatrix} \right\|_2.$$

It can easily verified that this least squares problem is equivalent to the original problem. But solving such a least square problem directly requires solving a system of equations which is M times larger in size.

4. The coefficient matrix \tilde{A} is independent of m , therefore, solution methods for multiple right hand side can be explored.
5. Some savings in computational cost can be achieved by solving (2.9) for $m = 1, 2, \dots, M - 1$ followed by setting $\alpha^M = \mathbf{1} - \sum_{m=1}^{M-1} \alpha^m$.

3. Properties of the Coefficient Matrix \tilde{A} . In this section, we study the properties of the coefficient matrix \tilde{A} in (2.9) and (2.10).

3.1. Discrete Maximum Principle. We first show that the solution of $\tilde{A}\alpha^m = \mathbf{b}^m$ satisfies a discrete maximum principle which implies \tilde{A} is non-singular and the bilateral constraints $\mathbf{0} \leq \alpha^m \leq \mathbf{1}$ are automatically satisfied.

We first recall the (strong) discrete maximum principle. Let Ω be a discrete connected domain and let Γ be its boundary ($\Omega \cap \Gamma = \emptyset$). Let $\alpha = (\alpha_i)$ be a discrete

function defined on $\Omega \cup \Gamma$. Then the principle says that α attains its maximum in Γ only, unless α is a constant everywhere in $\Omega \cup \Gamma$. The minimum principle can be defined similarly. The connectedness of Ω is defined base on the matting neighbor N_i . Pixel i and pixel j are connected if and only if they are a neighbor of each other. That is, w_{ij} and w_{ji} are both positive. If Ω has more than one connected component, then the principle can be applied to each component independently. Therefore, we can safely assume that Ω is connected. As we mentioned in Section 2.3, we also assume that each point in Γ is a matting neighbor of some pixel in Ω .

Next, we show that the optimal alpha mattes satisfy the maximum principle. The minimum principle can be proved similarly.

THEOREM 3.1. *The solution to $\tilde{A}\alpha = \mathbf{b}$ satisfies the discrete maximum principle. Here, $\mathbf{b} = (b_i)$ is a vector of the form*

$$\mathbf{b} = \begin{bmatrix} \mathbf{0} \\ \mathbf{b}_\Gamma \end{bmatrix}.$$

Proof. We assume that the maximum is attained at an interior point $i_0 \in \Omega$. Then, $(\tilde{A}\alpha)_{i_0} = 0$ because $b_{i_0} = 0$. Moreover, the i_0 -th row of \tilde{A} is the same as the i_0 -th row of $A = I - W$. Thus,

$$(\tilde{A}\alpha)_{i_0} = \alpha_{i_0} - \sum_{j \in N_{i_0}} w_{i_0j} \alpha_j = 0$$

or

$$\alpha_{i_0} = \sum_{j \in N_{i_0}} w_{i_0j} \alpha_j.$$

Since $w_{i_0j} > 0$ for $j \in N_{i_0}$ and $\sum_{j \in N_{i_0}} w_{i_0j} = 1$, the above equation says that the maximum value α_{i_0} equals to a weighted average of $\{\alpha_j : j \in N_{i_0}\}$. This implies that $\alpha_{i_0} = \alpha_j$ for all $j \in N_{i_0}$. For each $j \in N_{i_0}$, α_j is also the maximum so that the same argument can be applied to conclude that the value of α on N_j is again the maximum. Since the domain Ω is connected, we can conclude that α attains its maximum at every point in Ω and in the matting neighbor of Ω which covers Γ . This shows that if α has an interior maximum, then α is constant everywhere in $\Omega \cup \Gamma$. \square

In particular, the maximum (and minimum) principle implies that the strict bilateral constraints $\mathbf{0} < \alpha < \mathbf{1}$ and uniqueness of solution. We present these results in the next two corollaries.

COROLLARY 3.2. *The solution to $\tilde{A}\alpha = \mathbf{b}$ automatically satisfies the strict bilateral constraints*

$$\mathbf{0} < \alpha < \mathbf{1}$$

provided that $\Gamma = \{i \in \Gamma : b_i = 0\} \cup \{i \in \Gamma : b_i = 1\}$ and both sets are non-empty.

Proof. By the maximum principle, we $\alpha \leq \max_{i \in \Gamma} \alpha_i = 1$ and $\alpha \geq \min_{i \in \Gamma} \alpha_i = 0$. Since α is not constant on Γ , it cannot have any interior maximum nor interior minimum. Therefore, the bilateral constraints are strict in Ω . \square

COROLLARY 3.3. *The solution to $\tilde{A}\alpha = \mathbf{b}$ is unique.*

Proof. Let $\mathbf{b} = \mathbf{0}$, i.e. all boundary pixels have zero alpha matte. Then, by the maximum principle, we have $\alpha_\Omega \leq \max_{i \in \Gamma} \alpha_i = 0$ and $\alpha_\Omega \geq \min_{i \in \Gamma} \alpha_i = 0$. Thus,

$\alpha = \mathbf{0}$. This shows that the only solution to $\tilde{A}\alpha = \mathbf{0}$ is $\alpha = \mathbf{0}$. Therefore, \tilde{A} is non-singular. \square

A direct consequence is that the submatrix $A_{\Omega\Omega}$ of \tilde{A} is also non-singular, a fact we used when deriving the optimality conditions (2.7)–(2.8).

COROLLARY 3.4. *The matrix $A_{\Omega\Omega}$ is non-singular.*

Proof. We note that \tilde{A} is block upper triangular with diagonal blocks $A_{\Omega\Omega}$ and I . Thus, $0 \neq \det(\tilde{A}) = \det(A_{\Omega\Omega})\det(I) = \det(A_{\Omega\Omega})$ where $\det(\cdot)$ denotes the determinant. \square

3.2. M-matrix. Next we show that \tilde{A} is an M -matrix. This property will be useful when we design some effective preconditioners for \tilde{A} . We recall that a matrix A is called an M -matrix if

1. $a_{ii} > 0$ for all i ;
2. $a_{ij} \leq 0$ for all $i \neq j$;
3. A is non-singular;
4. $A^{-1} \geq 0$ (entry-wise).

It can be shown that the third and the fourth conditions can be replaced by $\rho(I - D^{-1}A) < 1$ where $\rho(\cdot)$ denotes the spectral radius (the largest eigenvalue in absolute value) and D is the diagonal part of A , see [9, p.29].

THEOREM 3.5. *The coefficient matrix \tilde{A} is an M -matrix.*

Proof. It is obvious that $\tilde{A} = (\tilde{a}_{ij})$ satisfies $\tilde{a}_{ii} = 1 > 0$ for all i and $\tilde{a}_{ij} = -w_{ij} \leq 0$ for all $i \neq j$. The diagonal part of \tilde{A} is I , and therefore, it remains to show that $\rho(I - \tilde{A}) < 1$.

We partition W into two submatrices:

$$W = \begin{bmatrix} W_{\Omega} \\ W_{\Gamma} \end{bmatrix}.$$

Then, we have

$$\tilde{W} := I - \tilde{A} = \begin{bmatrix} W_{\Omega} \\ \mathbf{0} \end{bmatrix}.$$

We note that the row sum $\tilde{W}\mathbf{1}$ is a binary vector. Moreover, $\tilde{W} \geq 0$ so that $\|\tilde{W}\|_{\infty} = \|\tilde{W}\mathbf{1}\|_{\infty} = 1$. Since $\rho(\cdot) \leq \|\cdot\|$ for any matrix norm, $\rho(\tilde{W}) \leq \|\tilde{W}\|_{\infty} = 1$. To show the strictly inequality, we assume to the contrary that $\rho(\tilde{W}) = 1$. Then \tilde{W} has an eigenvector $\mathbf{x} \neq \mathbf{0}$ and a corresponding eigenvalue ν (complex-valued) with $|\nu| = 1$. So, $\tilde{W}\mathbf{x} = \nu\mathbf{x}$. In particular, $|(\tilde{W}\mathbf{x})_i| = |\nu||x_i| = |x_i|$. If $i \in \Omega$, then $|\sum_j w_{ij}x_j| = |x_i|$ so that the arguments used to prove the maximum principle can be applied to argue that $|x_i| \equiv \text{constant}$ for all $i \in \Omega \cup \Gamma$. If $i \in \Gamma$, then we have $(\tilde{W}\mathbf{x})_i = 0$ so that $x_i = 0$. Together with the fact that $|x_i|$ is constant, we have $x_i = 0$ for all $i \in \Omega \cup \Gamma$. This contradicts to $\mathbf{x} \neq \mathbf{0}$. Thus, $\rho(\tilde{W}) < 1$. \square

4. Numerical Solvers. We have shown that the multiple-class image matting problem amounts to solving systems of linear equations of the form $\tilde{A}\alpha = \mathbf{b}$. Typically, we are dealing with a very large system as the number of pixels can be very large. We note that \tilde{A} is a sparse matrix because the number of nonzeros in each row depends on the size of geometric and photometric neighbors. Since the nonzero positions are image dependent (i.e., \tilde{A} is not fixed), the computer may run out of memory due to the problem of fill-in for some nonzero patterns of \tilde{A} . Direct methods may not be effective. As \tilde{A} is sparse, its matrix-vector multiplication can be computed very efficiently.

Thus we apply iterative methods [9] to solving linear systems arising from the optimality conditions of optimization problems. In particular, we study domain decomposition (DD) methods and the preconditioned GMRES method with domain decomposition preconditioners for solving such systems. We demonstrate that they are quite effective and outperform the other tested iterative methods.

In the domain decomposition method, we first decompose the image domain $\Omega \cup \Gamma$ into possibly overlapping regions $\Omega_1, \Omega_2, \dots, \Omega_L$. Let $\alpha^{(k)}$ be the k -th iterate. For each $l = 1, 2, \dots, L$, we solve the smaller system

$$\tilde{A}_{\Omega_l \Omega_l} \alpha_{\Omega_l}^{(k+1/2)} = \mathbf{b}_{\Omega_l} - \tilde{A}_{\Omega_l \Omega'_l} \alpha_{\Omega'_l}^{(k)}$$

for $\alpha_{\Omega_l}^{(k+1/2)}$. Here Ω'_l is the complement of Ω_l and \tilde{A} is partitioned into

$$\tilde{A} = \begin{bmatrix} \tilde{A}_{\Omega_l \Omega_l} & \tilde{A}_{\Omega_l \Omega'_l} \\ \tilde{A}_{\Omega'_l \Omega_l} & \tilde{A}_{\Omega'_l \Omega'_l} \end{bmatrix}.$$

Then, we update the new iterate by $\alpha_{\Omega_l}^{(k+1)} = \alpha_{\Omega_l}^{(k+1/2)}$ for each l . If pixel i belongs to more than one region, then the value of $\alpha_i^{(k+1)}$ is defined to be the average of the contributions from the several regions containing i . This method is also known as the block Jacobi method [9]. Similarly, the block Gauss Seidel method can be considered.

Next, we state that the iterations are well-defined because the matrix $\tilde{A}_{\Omega_l \Omega_l}$ is non-singular.

THEOREM 4.1. *The matrix $\tilde{A}_{\Omega_l \Omega_l}$ is non-singular for each l .*

Proof. Let B be the matrix

$$B = \begin{bmatrix} \tilde{A}_{\Omega_l \Omega_l} & 0 \\ 0 & I \end{bmatrix}.$$

Then we have $\tilde{A} \leq B$ and $b_{ij} \leq 0$ for all $i \neq j$. Recall that \tilde{A} is an M -matrix. Then, by Theorem 1.18 in [9], B is also an M -matrix. In particular, B , and hence $\tilde{A}_{\Omega_l \Omega_l}$, is non-singular. \square

In the implementation, these smaller systems can be solved by direct methods. Moreover, since the matrices $\tilde{A}_{\Omega_l \Omega_l}$ are fixed over the iterations, we can compute and store their LU factorization once and for all. In general, the regions Ω_l 's should be chosen to be as large as possible to encourage a fast convergence. But they should also be small enough so that direct methods can be used efficiently. On the other hand, the convergence rate generally increases with the amount of overlapping between regions. A good feature of this method is that it can be applied to very huge problems. We just have to decompose the domain so that Ω_l 's are small enough. The downside of having too many small regions is a slower convergence. Nevertheless the memory problem is solved. Moreover, this method can be easily parallelized.

5. Experimental Results.

5.1. Efficiency of Numerical Solvers. In this subsection, we test the performance of the numerical solvers discussed in §4. Two test images are shown in Figures 5.1 and 5.2 which are used to solve a two-class and a three-class matting problems respectively. In the figures, we show an original image, juxtaposed with the boundary conditions (scribbles) put on by the user. We also show the resulting alpha mattes and synthesized images using the mattes. We observe that the tiger and

the background are nicely separated in Figure 5.1 whereas the trees, the sky and the sea are also gracefully separated in Figure 5.2. Moreover, the alpha mattes allow a smoother transition from one segment to another than hard segmentation. Therefore, the synthesized images look natural.




FIG. 5.1. *Two-class image matting. (a) The original 384×256 Tiger image juxtaposed with initial scribbles put on by the user indicating foreground (red) and background (blue). (b) The alpha matte pulled. (c) The foreground viewed against a uniform blue background.*

In Figures 5.3, we compare the convergence of the relative residuals of point Jacobi method, point Gauss-Seidel method and several domain decomposition methods using the Tiger image in Figure 5.1(a) and the Sunset image in Figure 5.2(a). Domain decomposition methods include block Jacobi iterations with or without overlap and block Gauss-Seidel iterations with or without overlap. We find that the point Jacobi method, point Gauss-Seidel method and domain decomposition methods without overlap cannot reach a high level of accuracy within 100 iterations. Indeed, it essentially stagnates after 100 iterations. On the other hand, the convergence of domain decomposition methods with overlap is significantly faster.

Next, we study the effect of the amount of overlapping in the domain decomposition method. In Figure 5.4(a), we show the results using the Tiger image. Here, the domain is decomposed into 3×2 blocks where the size of each block is varied so that the number of overlapping pixels with its adjacent block is 0, 8, 16 and 32 pixels. In Figure 5.4(b), we show the results using the Sunset image. Here, the domain is decomposed into 3×4 blocks where the size of each block is varied so that the number of overlapping pixels with its adjacent block is 0, 8, 16 and 32 pixels. The iterations are stopped if the relative residual drops below 10^{-6} . We observe that the larger the number of overlapping pixels, the faster the convergence rate. Therefore, in practice, one may want to use as many overlapping pixels as possible while keeping the size of the linear systems to solve in each DD iteration manageable.

We can employ domain decomposition methods as preconditioners for the preconditioned GMRES method. In Figures 5.5, we compare the convergence of the relative residuals of different preconditioners: point Jacobi method, point Gauss-Seidel method and domain decomposition preconditioners of block Jacobi iterations with or without overlap and block Gauss-Seidel iterations with or without overlap.


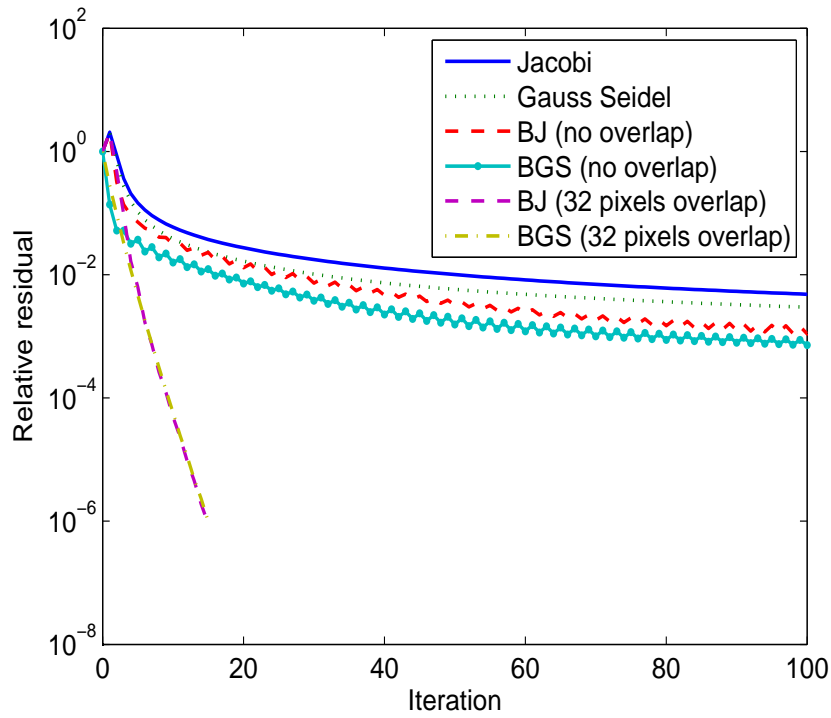
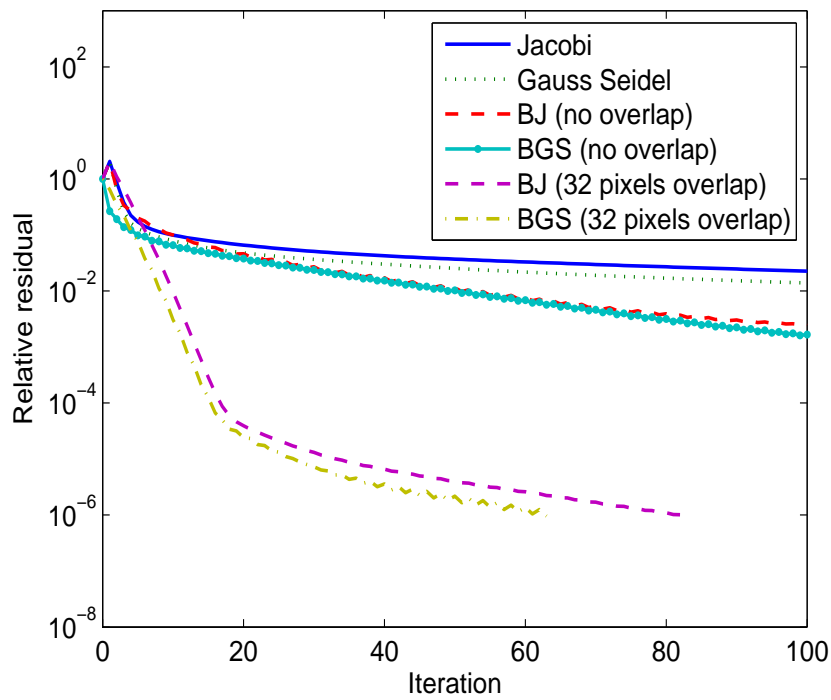


FIG. 5.2. Three-class image matting. (a) The original 480×640 Sunset image juxtaposed with initial scribbles put on by the user indicating three regions (blue, red and green). (b), (d), (f) The alpha matte for the blue, red and green region pulled respectively. (c), (e), (g) The blue, red and green region viewed against a uniform blue background respectively.




(a)

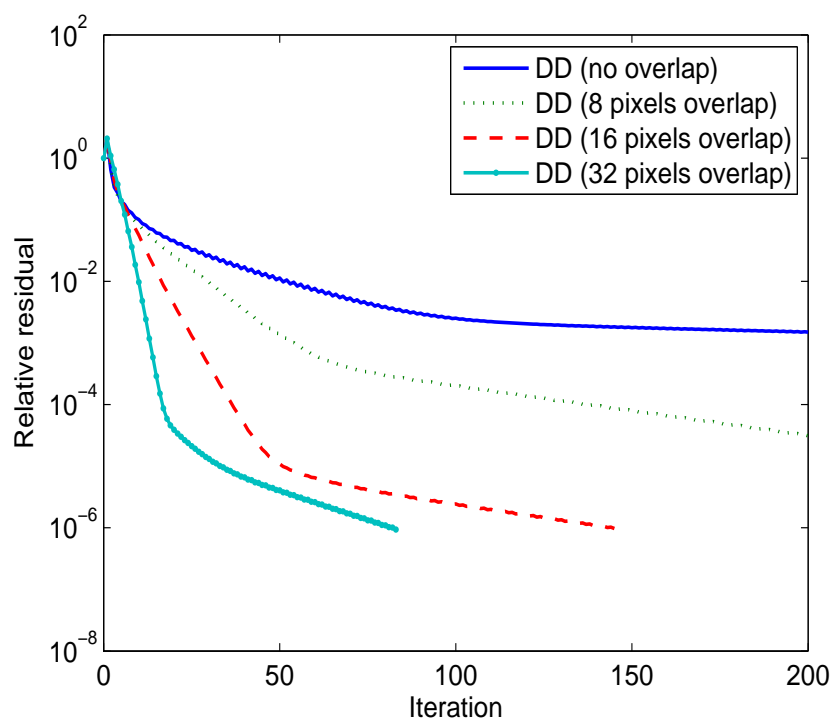


(b)

FIG. 5.3. Convergence profile of various iterative methods. (a) The Tiger image. (b) The Sunset image.



(a)



(b)

FIG. 5.4. Convergence profile of domain decomposition method with various degree of overlapping. (a) The Tiger image. (b) The Sunset image.

We find that the convergence behavior of the preconditioned GMRES method with all the tested preconditioners are improved with respect to the convergence results in Figure 5.3. In particular, the convergence of the preconditioned GMRES method with domain decomposition preconditioners of block Jacobi and block Gauss-Seidel iterations with overlap is very quick.

5.2. Effect of Boundary Conditions. In the image matting model, the user needs to specify some boundary conditions for each class of pixels. Therefore, it is useful to study how the boundary conditions should be specified. In this subsection, we carry a numerical study of the effect of boundary condition specification. We use a synthetic image and a real image, shown in Figures 5.6(a) and 5.7(a) respectively.

5.2.1. Location of boundary conditions. We first segment each image into foreground and background. In Figure 5.6(a), the foreground is the white square in the middle; the background is the black rim around the foreground. In Figure 5.7(a), the segmentation into foreground and background is done manually and is shown in Figure 5.7(b). Next, we use the pixel at the center of the image as the boundary condition for the foreground. One pixel in the background segment is used as the boundary condition for the background. This experiment is repeated using all possible locations of the background boundary condition.


To assess the effect of the location of the boundary condition, we compute the condition number of the corresponding coefficient matrix A . A smaller condition number is more desirable because the linear system can in general be solved faster. Intuitively, when boundary conditions are well specified, it should be easy to tell whether each pixel belongs the foreground or the background. This is in turn reflected in the condition number of A .

In Figures 5.6(d) and 5.7(e), the plots of the logarithmic condition number versus the background boundary condition location are shown. We see that the condition number increases significantly as the location is getting close to the foreground. This is expected because the feature vectors are computed over a small window. Therefore, when the background boundary condition is close to the foreground, the feature vector is computed using pixels from both the background and foreground. As a result, such a feature vector can be very different from those deep inside the background and those deep inside the foreground. This makes the classification of the true background pixels difficult because their feature vectors do not match any of the feature vectors from the boundary conditions. Another interesting observation from Figure 5.6(d) is that the condition number is smaller at the middle of the rim and is getting larger as we move towards either the inner or the outer boundary of the background segment.


In Figures 5.6(b) and 5.7(c), the alpha mattes resulted from the location leading to the best conditioned A are shown. The result in Figure 5.6(b) is very reasonable. But the result in Figure 5.7(c) is not ideal. It is because the picture in Figure 5.7(a) is quite complex, and therefore, boundary conditions with more number of pixels are needed to obtain good results.

In Figures 5.6(c) and 5.7(d), the alpha mattes resulted from the location leading to the worst conditioned A are shown. We can observe that these alpha mattes are of much poorer quality.

5.2.2. Size of boundary conditions. In this test, we set out to investigate the effect of the number of pixels used in the boundary conditions. We randomly pick $k = 5, 10, 15, 20, \dots, 800$ pixels from the image to be used as boundary conditions and record the condition number of the corresponding coefficient matrix A . Note that



(a)



(b)

FIG. 5.5. Convergence profile of the preconditioned GMRES method with domain decomposition preconditioners. (a) The Tiger image. (b) The Sunset image.

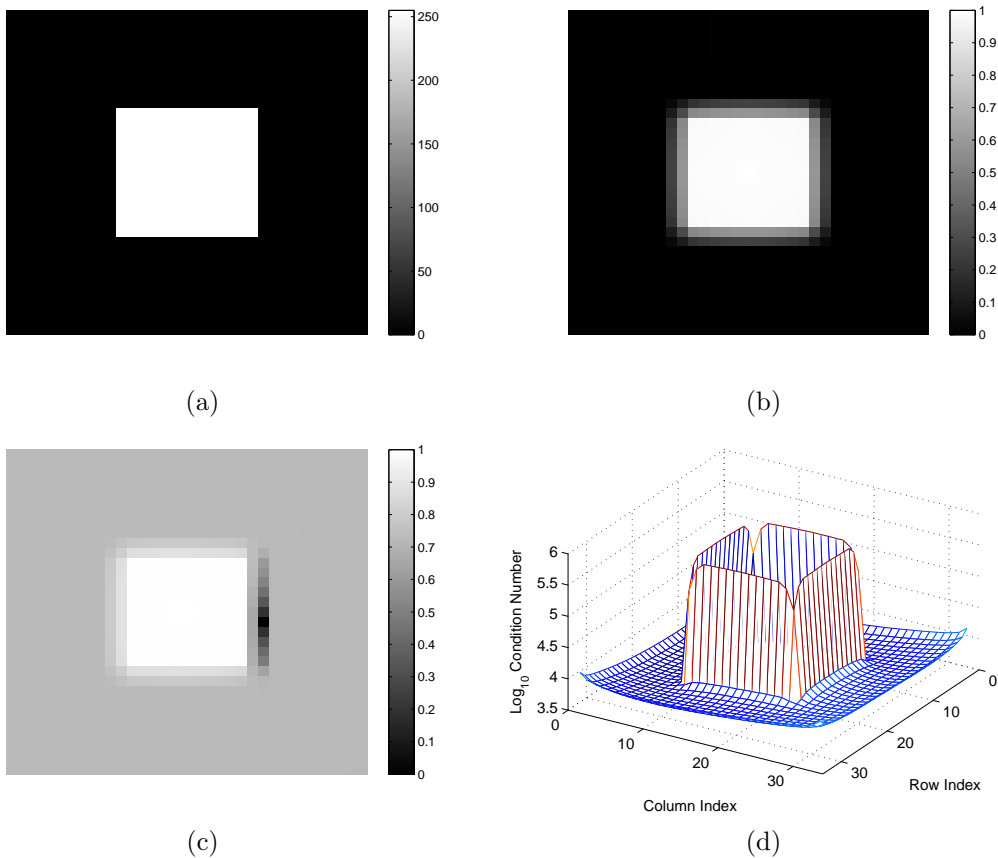


FIG. 5.6. *Effect of location of boundary conditions (synthetic image). (a) The input 33×33 image with a black background and a white foreground. The foreground boundary condition is given at the pixel at the center of the image. (b) The alpha matte obtained in the best conditioned case, where the background boundary condition is given at $(10, 6)$. (c) The alpha matte obtained in the worst conditioned case, where the background boundary condition is given at $(18, 24)$. (d) Logarithmic condition number of A with respect to location of the background boundary condition.*

the matrix A is independent of whether each chosen pixel is used as foreground or background boundary condition. For each fixed number of pixels k , we repeat the test 10 times with different sets of k pixels chosen. Then, the average condition number and the average \pm one standard deviation are reported. The resulting plots for the Square image in Figure 5.6(a) and the Black Swallowtail image in Figure 5.7(a) are shown in Figures 5.8(a) and (b) respectively. As expected, the average condition number decreases as more pixels as specified as boundary conditions. Moreover, the decrease is more significant initially.

6. Concluding Remarks. In this paper, we have studied theoretically and numerically the problem of multiple class image segmentation. We show that the existence and uniqueness of the optimal alpha matte. We also show through the discrete maximum principle that the bilateral constraints $0 \leq \alpha \leq 1$ are automatically satisfied so that the simpler unconstrained counterpart can be used instead. We propose to use domain decomposition methods and the preconditioned GMRES method


with domain decomposition preconditioners to solve the optimality condition. Our experimental results show that domain decomposition preconditioners with overlap are very effective for such linear systems. Furthermore, we have studied numerically the effect of boundary conditions and shown that the condition number of A can in general be decreased by specifying the boundary conditions at the central location of the respective region and by using more pixels as boundary conditions.

REFERENCES


- [1] Y. Boykov and M. Jolly, Interactive graph cuts for optimal boundary and region segmentation of objects in N-D images, Proc. IEEE Int. Conf. on Computer Vision, 2001.
- [2] Y. Boykov and V. Kolmogorov, An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision, IEEE Trans. PAMI, V26 (2004), pp. 1124-1137.
- [3] C. Rother, V. Kolmogorov and A. Blake, "GrabCut": Interactive foreground extraction using iterated graph cuts, ACM Trans. Graph., V23(3) (2004), pp. 309-314.
- [4] J. Wang and M. Cohen, An iterative optimization approach for unified image segmentation and matting, Proc. IEEE Int. Conf. on Computer Vision, 2005.
- [5] S. Yu and J. Shi, Segmentation given partial grouping constraints, IEEE Trans. PAMI, V26 (2004), pp. 173-183.
- [6] J. Guan and G. Qiu, Interactive image matting using optimization, technical report 01-2006, School of Computer Science and Information Technology, University of Nottingham, U.K., 2006.
- [7] J. Guan and G. Qiu, Interactive image segmentation using optimization with statistical priors, Proc. ECCV 2006, Graz, Austria, 2006.
- [8] A. Levin, D. Lischinski and Y. Weiss, Colorization using optimization, ACM Trans. Graph., V23 (2004), pp. 689-694.
- [9] Y. Saad, Iterative Methods for Sparse Linear Systems, PWS Publishing Company, 1996.
- [10] J. Shi and J. Malik, Normalized cuts and image segmentation, IEEE Trans. PAMI, V22 (2000), pp. 888-905.




(a)



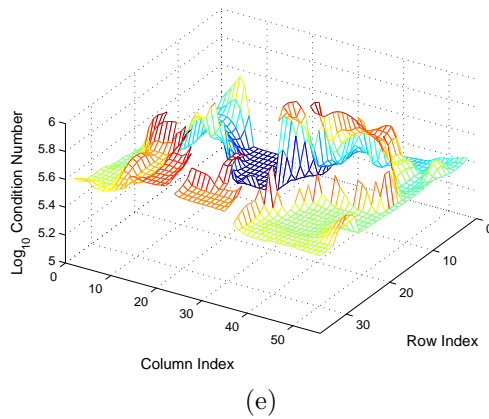
(b)



(c)




(d)

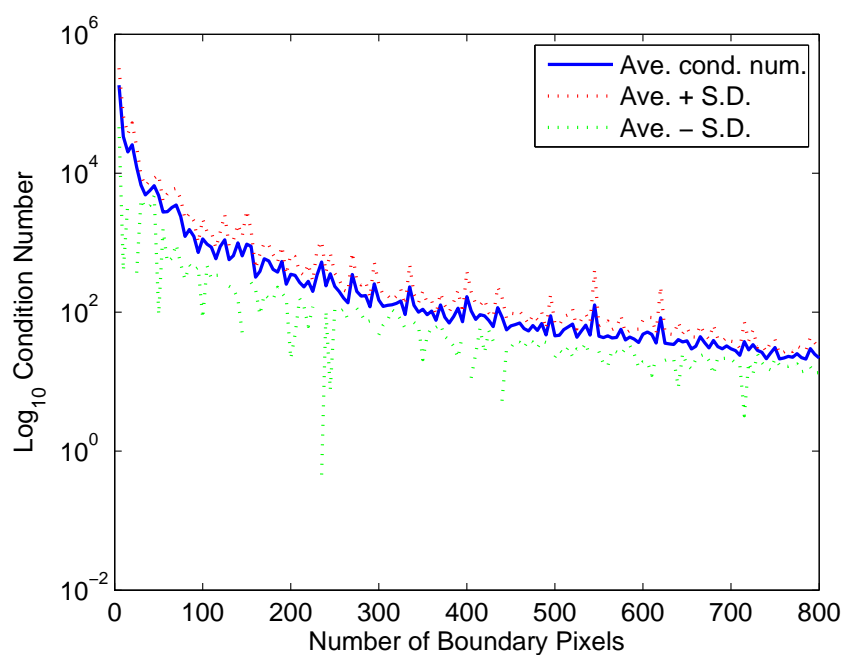


(e)

FIG. 5.7. Effect of location of boundary conditions (real image). (a) The input 35×55 Black Swallowtail image. (b) Foreground and background segmentation obtained manually. The foreground boundary condition is given at the pixel at the center of the image. (c) The alpha matte obtained in the best conditioned case, where the background boundary condition is given at pixel location (5, 21). (d) The alpha matte obtained in the worst conditioned case, where the background boundary condition is given at pixel location (25, 17). (e) Logarithmic condition number of A with respect to location of the background boundary condition.



(a)



(b)

FIG. 5.8. Effect of number of pixels used in the boundary conditions. The blue curve shows the average logarithmic condition number of A versus the number of randomly chosen boundary pixels. The red (green) curve shows the average plus (minus) one standard deviation of the logarithmic condition number of A . (a) Results using the Square image in Figure 5.6(a). (b) Results using the Black Swallowtail image in Figure 5.7(a).