

# Spectral Matting

Anat Levin, Alex Rav-Acha, and Dani Lischinski

**Abstract**—We present *spectral matting*: a new approach to natural image matting that automatically computes a basis set of fuzzy *matting components* from the smallest eigenvectors of a suitably defined Laplacian matrix. Thus, our approach extends spectral segmentation techniques, whose goal is to extract hard segments, to the extraction of soft matting components. These components may then be used as building blocks to easily construct semantically meaningful foreground mattes, either in an unsupervised fashion or based on a small amount of user input.

**Index Terms**—Image matting, unsupervised segmentation, spectral analysis.

## 1 INTRODUCTION

DIGITAL matting is the process of extracting a foreground object from an image along with an opacity estimate for each pixel covered by the object. This operation enables compositing the extracted object over a novel background and thus constitutes an invaluable tool in image editing, video production, and special effects in motion pictures.

In particular, the challenging case of *natural* image matting, which poses no restrictions on the background, has received much research attention. Recognizing that the problem is inherently underconstrained, all of the existing methods require the user to provide additional constraints in the form of a trimap [2], [22], [8] or a set of brush strokes [24], [14], [9] (while brush strokes are also a form of trimap, the user input is typically sparser). The question of whether (or to what degree) it is possible to automate the matting process is of considerable theoretical and practical interest.

In this paper, we attempt to provide some new insights into this question. Our work is strongly influenced by spectral segmentation methods [20], [25], [17], [26]. These methods analyze the smallest eigenvectors of the image's graph Laplacian matrix in order to obtain an unsupervised decomposition of the image into a collection of hard segments. In this work, we extend this idea to unsupervised computation of a collection of soft *matting components*.

Spectral segmentation methods, such as [20], resort to computation of real-valued eigenvectors as an approximation necessary to transform an NP-complete optimization problem into a tractable one. In contrast, we are not seeking a disjoint image partitioning, but rather attempting to

recover the fractional foreground coverage at each pixel. Specifically, we obtain our real-valued matting components via a linear transformation of the smallest eigenvectors of the *matting Laplacian* matrix, introduced by Levin et al. [14]. Once obtained, these matting components serve as building blocks for construction of complete foreground mattes.

This concept is illustrated in Fig. 1. Given the input image in Fig. 1a, one can produce an unsupervised disjoint hard partitioning of the image using, e.g., [26] (Fig. 1b). In contrast, we compute a set of overlapping, fractional, matting components, visualized in Fig. 1d. Combining three of these components (framed in red) yields the foreground matte of the girl, shown in Fig. 1c.

In summary, our main contribution is the introduction of the concept of fuzzy matting components and a method for computing them in an unsupervised manner. We then proceed to describe an unsupervised matting algorithm. Of course, just like unsupervised segmentation, unsupervised matting is an ill-posed problem. Thus, we focus more on two extensions that use our matting components to construct a particular matte: 1) present the user with several matting alternatives to choose from or 2) let the user specify her intent by just a few mouse clicks.

A shorter version of this paper appeared in [15].

## 2 MATTING COMPONENTS

Matting algorithms typically assume that each pixel  $I_i$  in an input image is a linear combination of a foreground color  $F_i$  and a background color  $B_i$ :

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

This is known as the *compositing equation*. In this work, we generalize the compositing equation by assuming that each pixel is a convex combination of  $K$  image layers  $F^1, \dots, F^K$ :

$$I_i = \sum_{k=1}^K \alpha_i^k F_i^k. \quad (2)$$

The  $K$  vectors  $\alpha^k$  are the *matting components* of the image, which specify the fractional contribution of each layer to the final color observed at each pixel. The matting components are nonnegative and sum to one at every

- A. Levin is with MIT CSAIL, The Stata Center 32-D466, 32 Vassar Street, Cambridge, MA 02139. E-mail: alevin@csail.mit.edu.
- A. Rav-Acha is with the Department of Computer Science and Applied Math, Ziskind Building, room 33, The Weizmann Institute of Science, Rehovot, 76100 Israel. E-mail: reavach@gmail.com.
- D. Lischinski is with the School of Computer Science and Engineering, Room 73, Ross Building, The Hebrew University of Jerusalem, Givat Ram, Jerusalem 91904, Israel. E-mail: danix@cs.huji.ac.il.

Manuscript received 3 Nov. 2007; revised 17 Apr. 2008; accepted 4 June 2008; published online 13 June 2008.

Recommended for acceptance by S. Baker, J. Matas, and R. Zabih.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number TPAMISI-2007-11-0746.

Digital Object Identifier no. 10.1109/TPAMI.2008.168.




Fig. 1. Spectral segmentation and spectral matting. (a) Input image. (b) Hard segmentation. (c) Alpha matte. (d) Matting components computed by our method.

pixel. The intuitive motivation for having these components is that, similarly to the individual low-level fragments in an oversegmented image, they may be used to construct higher level, semantically meaningful foreground mattes, as demonstrated in Fig. 1.

A desirable, although not required, property of the matting components is *sparsity*: Each component should be either completely opaque or completely transparent over as many image pixels as possible. This means that areas of transition between the different layers are limited to a small number of pixels and each pixel is influenced by a small number of layers.

In this paper, we explore the relationship between the matting components and the eigenvectors of the matting Laplacian matrix [14]. Specifically, we show that, under certain assumptions, the matting components are spanned by the smallest eigenvectors of the matting Laplacian. We then propose a method for computing the matting components by finding an appropriate linear transformation and applying it to these eigenvectors.

### 3 SPECTRAL ANALYSIS

We start by briefly reviewing the basic theory of spectral segmentation methods [5], [6], [7], [4], [11], [18], [10], [20], [25], [17], [26]. These methods typically associate with the image an  $N \times N$  affinity matrix  $A$ , such as  $A(i, j) = e^{-d_{ij}/\sigma^2}$ , where  $d_{ij}$  is some measure of the distance between the

pixels (such as color difference and geometric distance). One can then define the Laplacian matrix  $L = D - A$ , where  $D$  is the diagonal matrix  $D(i, i) = \sum_j A(i, j)$ .  $L$  is a symmetric positive semidefinite matrix whose eigenvectors capture much of the image structure.<sup>1</sup>

Consider the ideal case where the affinity matrix  $A$  captures exactly the fact that an image is composed from several distinct clusters, or *connected components*. That is, a subset  $C$  of the image pixels is a connected component of the image if  $A(i, j) = 0$  for every  $i, j$  such that  $i \in C, j \notin C$ , and there is no subset of  $C$  which satisfies this property. Let  $m^C$  denote the indicator vector of the component  $C$ :

$$m_i^C = \begin{cases} 1 & i \in C \\ 0 & i \notin C \end{cases}$$

then  $m^C$  is a 0-eigenvector of  $L$  (i.e., an eigenvector with eigenvalue 0).

Now, suppose that the image consists of  $K$  connected components  $C_1, \dots, C_K$  such that  $\{1, \dots, N\} = \bigcup_{k=1}^K C_k$ , where  $C_k$  are disjoint subsets of pixels. In this case, the indicator vectors  $m^{C_1}, \dots, m^{C_K}$  are all independent, orthogonal 0-eigenvectors of  $L$ . However, computing the

1. In fact, most spectral segmentation papers consider normalized affinity matrices such as  $D^{-1}L$  or  $D^{-1/2}LD^{-1/2}$ . However, in this work, we focus on  $L$  itself as it is not clear how to justify the normalization in the case of the matting Laplacian. The problem is that the off-diagonal elements of the matting Laplacian can be both negative and positive and, thus, the matting cost cannot be expressed as a sum of positive pairwise terms.

eigenvectors of  $L$  yields these indicator vectors only up to rotation. This is the case since, for any  $K \times K$  rotation matrix  $R$ , the vectors  $[m^{C_1}, \dots, m^{C_K}]R$  are also a basis for the nullspace of  $L$ .

In real images, the affinity matrix  $A$  is rarely able to perfectly separate between the different pixel clusters. Therefore, the Laplacian  $L$  usually does not have multiple 0-eigenvectors. However, it has been observed that the smallest eigenvectors of  $L$  tend to be nearly constant within coherent image components. Extracting the different components from the smallest eigenvectors is known as *spectral rounding* and has attracted much attention [17], [26], [23], [27], [13]. The simplest approach [17] is to cluster the image pixels using the  $k$ -means algorithm and use perturbation analysis to bound the error of this algorithm as a function of the connectivity within and between clusters. Other more recent methods [26], [27], which inspired the approach taken in this work, explicitly search for a rotation matrix that brings the eigenvectors as close as possible to binary indicator vectors.

### 3.1 Spectral Analysis with the Matting Laplacian

Our goal in this work is to derive an analogy between hard segmentation and matting and to show that fuzzy matting components may be extracted from the smallest eigenvectors of the matting Laplacian, similarly to the extraction of hard clusters described earlier.

#### 3.1.1 The Matting Laplacian

The matting Laplacian was introduced by Levin et al. [14] in order to evaluate the quality of a matte without explicitly estimating the foreground and background colors in (1). They show that if the colors of the background and the foreground within a local image window  $w$  form two different lines in RGB space, then the  $\alpha$  values within  $w$  may be expressed as a linear combination of the color channels:

$$\forall i \in w \quad \alpha_i = a^R I_i^R + a^G I_i^G + a^B I_i^B + b. \quad (3)$$

Thus, the matte extraction problem becomes one of finding the alpha matte that minimizes the deviation from the linear model (3) over all image windows  $w_q$ :

$$J(\alpha, a, b) = \sum_{q \in I} \sum_{i \in w_q} \left( \alpha_i - a_q^R I_i^R - a_q^G I_i^G - a_q^B I_i^B - b_q \right)^2 + \varepsilon \|a_q\|^2, \quad (4)$$

where  $\varepsilon \|a_q\|^2$  is a regularization term on  $a$ . The linear model coefficients  $a, b$  may be eliminated from (4), yielding a *quadratic cost in  $\alpha$  alone*:

$$J(\alpha) = \alpha^T L \alpha. \quad (5)$$

This cost has a trivial minimum, which is a constant  $\alpha$  vector, and thus, in the user-assisted framework described in [14],  $J(\alpha)$  is minimized subject to user constraints.

In (5),  $L$  is the *matting Laplacian*, a sparse symmetric positive semidefinite  $N \times N$  matrix whose entries are a function of the input image in local windows, depending neither on the unknown foreground and background colors nor on the linear model coefficients.  $L$  is defined as a sum

of matrices  $L = \sum_q A_q$ , each of which contains the affinities among pixels inside a local window  $w_q$ :

$$A_q(i, j) = \begin{cases} \delta_{ij} - \frac{1}{|w_q|} \\ \left( 1 + (I_i - \mu_q)^T \left( \Sigma_q + \frac{\varepsilon}{|w_q|} I_{3 \times 3} \right)^{-1} (I_j - \mu_q) \right) & (i, j) \in w_q \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

Here,  $\delta_{ij}$  is the Kronecker delta,  $\mu_q$  is the  $3 \times 1$  mean color vector in the window  $w_q$  around pixel  $q$ ,  $\Sigma_q$  is a  $3 \times 3$  covariance matrix in the same window,  $|w_q|$  is the number of pixels in the window, and  $I_{3 \times 3}$  is the  $3 \times 3$  identity matrix.

Note that the matting Laplacian is fundamentally different from a “standard” graph Laplacian matrix. Standard Laplacian matrices are constructed using non-negative affinities  $A(i, j)$  and, hence, all off-diagonal elements of  $L$  are nonpositive. On the other hand, according to (6), the off-diagonal elements of the matting Laplacian may have arbitrary signs. Thus, although we follow the terminology of Levin et al. [14] and refer to the matrix  $L$  as the matting Laplacian, it should be understood that this is an extension of the standard term as most important properties of graph Laplacians rely on the nonnegative affinity assumption.

Fortunately, as we show below, some of the useful properties of standard graph Laplacian matrices apply to the matting Laplacian as well. The first such property is that both types of Laplacians are positive semidefinite matrices. For the standard graph Laplacian, this property is trivially implied from the fact that the pairwise affinities  $A(i, j)$  are nonnegative; for every  $N$ -dimensional vector  $x$ , it holds that  $x^T L x = \sum_{i,j} A(i, j) (x_i - x_j)^2 \geq 0$ . The matting Laplacian is also positive semidefinite, but, instead of positive pairwise terms, it can be factored as a sum of positive terms consisting of the affinities in  $3 \times 3$  windows; by construction of the matting Laplacian (see [14] for the exact derivation), for every  $N$ -dimensional vector  $x$ , it holds that

$$x^T A_q x = \min_{a,b} \sum_{i \in w_q} \left( x_i - a_q^R I_i^R - a_q^G I_i^G - a_q^B I_i^B - b_q \right)^2 + \varepsilon \|a_q\|^2 \geq 0,$$

which implies  $x^T L x = \sum_q x^T A_q x \geq 0$ .

Another useful property of the matting Laplacian (6) is that its smallest eigenvectors appear to capture information about the fuzzy cluster assignments of pixels in the image, even before any user-specified constraints are taken into account. This observation was already made by Levin et al. [14]; however, they made no use of the eigenvectors beyond presenting them to the user as guides for scribble placement. In this work, we show that the smallest eigenvectors of the matting Laplacian span the individual matting components of the image.

#### 3.1.2 The Matting Laplacian's Nullspace

To gain some understanding, we begin by studying the ideal case. To justify the usage of spectral analysis to estimate matting components, our goal is to show that,



Fig. 2. The smallest eigenvectors of the matting Laplacian for the image in Fig. 1a. Linear combinations of these eigenvectors produced the matting components shown in Fig. 1d.

under reasonable conditions, the actual matting components belong to the nullspace of the matting Laplacian. We say that a matting component  $\alpha^k$  is *active* in a local image window  $w$  if there exists a pixel  $i \in w$  for which  $\alpha_i^k > 0$ . The following claim states the conditions on the local color distribution in each layer, under which  $L\alpha^k = \mathbf{0}$ . The severity of the conditions is related to the number of active layers in a local window. The least restricted case is when only one layer is active, in which the local color distribution can be arbitrary complex. The most restricted case is when a window contains three active layers (as in the case of a T-junction) and, for such windows, the theorem holds when each layer color is locally uniform.

**Claim 1.** Let  $\alpha^1, \dots, \alpha^K$  be the actual decomposition of the image  $I$  into  $k$  matting components. The vectors  $\alpha^1, \dots, \alpha^K$  lie in the nullspace of the matting Laplacian  $L$  (given by (6) with  $\varepsilon = 0$ ) if every local image window  $w$  satisfies one of the following conditions:

1. A single component  $\alpha^k$  is active within  $w$ .
2. Two components  $\alpha^{k_1}, \alpha^{k_2}$  are active within  $w$  and the colors of the corresponding layers  $F^{k_1}, F^{k_2}$  within  $w$  lie on two different lines in RGB space.
3. Three components  $\alpha^{k_1}, \alpha^{k_2}$  and  $\alpha^{k_3}$  are active within  $w$ , each layer  $F^{k_1}, F^{k_2}, F^{k_3}$  has a constant color within  $w$ , and the three colors are linearly independent.

**Proof.** The matting cost (5) measures the deviation between a matte and a linear function of the color channels over all local windows (4). Thus, in order to show that a matte component  $\alpha^k$  satisfies  $L\alpha^k = \mathbf{0}$ , it suffices to show that, for every local window  $w$ , there exist  $a^R, a^G, a^B, b$  such that:  $\alpha_i^k = a^R I_i^R + a^G I_i^G + a^B I_i^B + b, \forall i \in w$ . Below, we show this for each of the three window types.

**Case 1.** Since the matting components sum to 1 at every image pixel, the single active component  $\alpha^k$  must equal 1 within  $w$ . Thus, it is easily expressed as a linear function of the image by setting  $a^R = a^G = a^B = 0$  and  $b = 1$ .

**Case 2.** This case is equivalent to [14, Theorem 2].

**Case 3.** Assume wlog that  $k_1 = 1, k_2 = 2, k_3 = 3$ . Let  $\mathbf{F} = [F^1, F^2, F^3]$  be a  $3 \times 3$  matrix of the uniform layer colors and, for  $i \in w$ , let  $\alpha_i = [\alpha_i^1, \alpha_i^2, \alpha_i^3]^T$  be a  $3 \times 1$  vector of components values. We note that  $I_i = \mathbf{F}\alpha_i$ .

Since the three layer colors are linearly independent,  $\mathbf{F}$  is invertible and  $\alpha_i = \mathbf{F}^{-1}I_i$ , which implies that  $\alpha^k$  are linear functions of the image.  $\square$

As in the case of standard Laplacians, when the smallest eigenvectors of the matting Laplacian are computed, the result may be any linear combination of the different matting components and recovering the individual components is equivalent to linearly transforming the eigenvectors. It should be noted that, unlike hard segments, the matting components are not binary vectors and, thus, are not necessarily orthogonal. Hence, while the eigenvectors are orthogonal, the transformation from eigenvectors to matting components might be a general linear transformation and not a simple rotation.

To summarize, the main conclusion of the above discussion is that, whenever the matting components of an image satisfy the conditions of Claim 1, they may be expressed as a linear combination of the 0-eigenvectors of  $L$ .

In most real images, the assumptions of Claim 1 do not hold exactly and, thus, the matting Laplacian might not have multiple 0-eigenvectors. Yet, if the layers are sufficiently distinct, they are generally captured by the smallest eigenvectors of  $L$ . For example, Fig. 2 shows the smallest eigenvectors for a real image, all exhibiting the fuzzy layer boundaries. We have empirically observed that the matting components of real images are usually spanned quite well by the smallest eigenvectors of the matting Laplacian. Indeed, the components shown in Fig. 1d were obtained as linear combinations of the smallest eigenvectors.

### 3.2 From Eigenvectors to Matting Components

As explained above, recovering the matting components of the image is equivalent to finding a linear transformation of the eigenvectors. Recall that the matting components should sum to 1 at each image pixel and they should be near 0 or 1 for most image pixels since the majority of image pixels are usually opaque. Thus, we are looking for a linear transformation of the eigenvectors that would yield a set of nearly binary vectors.




Fig. 3. Sparsity score:  $\rho(\alpha) = |\alpha|^{0.9} + |1 - \alpha|^{0.9}$  favoring binary  $\alpha$  values.

More formally, let  $E = [e^1, \dots, e^K]$  be the  $N \times K$  matrix of eigenvectors. Our goal is then to find a set of  $K$  linear combination vectors  $y^k$  that minimize

$$\begin{aligned} & \sum_{i,k} |\alpha_i^k|^\gamma + |1 - \alpha_i^k|^\gamma, \quad \text{where } \alpha^k = Ey^k \\ & \text{subject to } \sum_k \alpha_i^k = 1. \end{aligned} \quad (7)$$

If  $0 < \gamma < 1$  is used (in our implementation  $\gamma = 0.9$ ), then  $|\alpha_i^k|^\gamma + |1 - \alpha_i^k|^\gamma$  is a robust score measuring the sparsity of a matting component (plotted in Fig. 3). Without the requirement  $\alpha^k = Ey^k$ , the sparsity term would be minimized by binary vectors, but, as the vectors  $\alpha^k$  are restricted to linear combinations of the eigenvectors, they must maintain the fuzzy layer boundaries. Although we do not explicitly constrain the  $\alpha$  values to be between 0 and 1, in practice, the resulting values tend to lie in this range due to the sparsity penalty.

The above cost is, of course, a nonconvex one and we optimize it iteratively using Newton's method [3], which amounts to a sequence of second-order approximations. Given a guess for the values of  $\alpha_i^k$ , we define  $u_i^k$  as the second derivative of  $|\alpha_i^k|^\gamma$ ,  $u_i^k \propto |\alpha_i^k|^{(\gamma-2)}$ , and, similarly,  $v_i^k \propto |1 - \alpha_i^k|^{(\gamma-2)}$ . The second-order approximation to (7) reads as minimizing

$$\begin{aligned} & \sum_{i,k} u_i^k |\alpha_i^k|^2 + v_i^k |1 - \alpha_i^k|^2, \quad \text{where } \alpha^k = Ey^k \\ & \text{subject to } \sum_k \alpha_i^k = 1. \end{aligned} \quad (8)$$

As this problem has now become one of quadratic optimizations under linear constraints, the optimal solution can be computed in closed form by inverting a  $K^2 \times K^2$  matrix. For that, we define  $U^k = \text{diag}(u^k)$  and  $V^k = \text{diag}(v^k)$ . We also define  $W^k = E^T(U^k + V^k)E$  and  $\mathbf{1}$  as a vector of ones. It may be shown that the  $K^2$  elements of  $y^1, \dots, y^K$  are the solution for

$$\begin{aligned} & \begin{bmatrix} \text{Id} & \text{Id} & \text{Id} & \dots & \text{Id} \\ 0 & W^2 + W^1 & W^1 & \dots & W^K \\ 0 & W^1 & W^3 + W^1 & \dots & W^K \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & W^1 & W^1 & \dots & W^K + W^1 \end{bmatrix} y \\ & = \begin{bmatrix} E\mathbf{1} \\ Ev^2 + Eu^1 \\ Ev^3 + Eu^1 \\ \vdots \\ Ev^K + Eu^1 \end{bmatrix}, \end{aligned} \quad (9)$$

where  $\text{Id}$  is the  $K \times K$  identity matrix. Given an initial guess, our algorithm iteratively solves a sequence of second-order optimization problems of the form of (8) and uses the solution of each iteration to reweight the following one. We note that the weights  $u_i^k$  (and  $v_i^k$ ) are higher when the current guess for  $\alpha_i^k$  is close to 0 (or 1) and are lower when the guess is farther away from these values. Thus, the effect of the reweighting step is to pull toward 0 (or 1) those  $\alpha$  entries for which the current guess is already close and to loosen the relation for pixels for which the guess is far anyway. For example, if the current guess for  $\alpha_i^k$  is close to 0, then  $u_i^k \gg v_i^k$  and, thus, the term  $u_i^k |\alpha_i^k|^2 + v_i^k |1 - \alpha_i^k|^2$  pulls the alpha component at this pixel toward 0 much more strongly than toward 1.

Since the cost (7) is not convex, the result of the Newton process strongly depends on the initialization. One useful way to initialize the process is to apply a  $k$ -means algorithm on the smallest eigenvectors of the matting Laplacian and project the indicator vectors of the resulting clusters onto the span of the eigenvectors  $E$ :

$$\alpha^k = EE^T m^{C^k}. \quad (10)$$

The effect of this projection step is to find a linear combination of eigenvectors that minimizes the squared distance to  $m^{C^k}$ . We note that the resulting matting components sum to one since

$$\sum_k \alpha^k = EE^T \left( \sum_k m^{C^k} \right) = EE^T \mathbf{1} = \mathbf{1}. \quad (11)$$

The last equality follows from the fact that the constant vector belongs to the nullspace of the matting Laplacian. As a result, projecting the hard clusters provides a legal solution for (7). We also note that, despite the fact that  $m^{C^k}$  are binary vectors, the projection typically features fuzzy matting boundaries. This is due to the fact that, as illustrated in Fig. 2, the smallest eigenvectors all maintain the fuzzy layer structure and, thus, simply do not suffice to span the hard segmentation boundaries.

In practice, we typically use a larger number of eigenvectors than the number of matting components to be recovered. Using more eigenvectors makes it possible to obtain sparser components. The reason is that more basis elements span a richer set of vectors (in the extreme case, if all  $N$  eigenvectors are used, any binary vector can be generated). A number of examples demonstrating the extraction of soft matting components are given in Figs. 1 and 4.




Fig. 4. A number of test images and the matting components extracted from them using our method. A random color is assigned to each component for visualization purposes.

## 4 GROUPING COMPONENTS

So far, we have shown how matting components may be extracted from the matting Laplacian. However, usually the matting components are not a goal in their own as one is ultimately interested in recovering a complete matte for some foreground object. Fortunately, all that is needed to obtain a complete matte is to specify which of the components belong to the foreground. Let  $b$  denote a  $K$ -dimensional binary vector indicating the foreground components (i.e.,  $b^k = 1$  iff  $\alpha^k$  is a foreground component). The complete foreground matte is then obtained by simply adding the foreground components together:

$$\alpha = \sum_k b^k \alpha^k. \quad (12)$$

For example, the matte in Fig. 1c was obtained by adding the components highlighted in red in Fig. 1d.

For the applications discussed below, one would like to compare multiple grouping hypotheses. If the smallest eigenvalues are not exactly zero (which is the case for most real images), we can also measure the quality of the resulting  $\alpha$ -matte as  $\alpha^T L \alpha$ , where  $L$  is the matting Laplacian (6). When a large number of hypotheses is to be tested, multiplying each hypothesis by  $L$  might be too expensive. However, if each hypothesis is just a sum of matting components, we can precompute the correlations between the matting components via  $L$  and store them in a  $K \times K$  matrix  $\Phi$ , where

$$\Phi(k, l) = \alpha^k{}^T L \alpha^l. \quad (13)$$

The matte cost can then be computed as

$$J(\alpha) = b^T \Phi b, \quad (14)$$

where  $b$  is a  $K$ -dimensional binary vector indicating the selected components. Thus, if  $\Phi$  has been precomputed,  $J(\alpha)$  can be evaluated in  $O(K^2)$  operations instead of  $O(N)$  operations.

### 4.1 Unsupervised Matting

Given an image and a set of matting components, we would like to split the components into foreground and background groups and pull out a foreground object. If the grouping criterion takes into account only low-level cues, then we just search for a grouping with the best matting cost, as defined in (14). However, the matting cost is usually biased toward mattes which assign nonconstant values only to a small subset of the image pixels (in the extreme case, the best matte is a constant one). The spectral segmentation literature suggests several criteria that overcome this bias. One approach is to search for quotient cuts (e.g., normalized cuts [20]) which score a cut as the ratio between the cost of the cut and the size of the resulting clusters. A second approach is to look for balanced cuts [13], where the size of each cluster is constrained to be above a certain percent of the image size. In this work, we follow this latter approach and rule out trivial solutions by considering only groupings that assign at least 30 percent of the pixels to the foreground and at least 30 percent of the pixels to the background. When the number  $K$  of matting components is small, we can enumerate all  $2^K$  hypotheses and select the one with the best score using (14).

Fig. 5 shows some results produced by the unsupervised matting approach described above. In two of these examples, the hypothesis with the highest score indeed corresponds to the “correct” foreground matte, while, in one example (bottom row), the “correct” matte was ranked fourth. Note that, in all of these examples, the other high-ranked hypotheses are quite sensible as well, considering that our approach does not attempt to perform any high-level image understanding. Of course, it is not hard to find examples where unsupervised matting fails, and the last two examples in Fig. 5 illustrate such failures. In general, whenever the foreground or background objects consist of several visually distinct components, the assignment with the minimal matting cost might not correspond to our visual perception. In fact, it is well known within the image segmentation community that, while unsupervised bottom-up cues can



Fig. 5. Unsupervised matting results for a few images. The hypotheses are ordered according to their score.

efficiently group coherent regions in an image, the general image segmentation problem is inherently ambiguous and requires additional information. In practice, such as in the case of hard image segmentation, the foreground/background assignment may be guided by several additional cues such as top-down models [1], color statistics [19], or motion and focus cues. Thus, we believe that the main practical use of matting components is in the supervised setting and focus on user-guided matting in the remainder of this paper.

## 4.2 User-Guided Matting

We now consider an interactive setting where the user guides the matting process toward the extraction of the desired foreground matte. In such a setting, the foreground/background assignment of some of the components is determined by the user, thereby reducing the number of legal hypotheses to be tested. Given very minimal foreground and background constraints, it is usually possible to rule out trivial solutions, so there is no need to explicitly keep the size of each group above a certain threshold (as in the unsupervised case). In this case, we can speed up the search for an optimal foreground/

background assignment of components using a graph min-cut formulation.

### 4.2.1 Optimizing Assignments in the Components Graph

To efficiently search for the optimal foreground/background assignment, we approximate the matting cost (14) as a sum of pairwise terms. This enables us to approximate the search for the optimal foreground/background assignment as a min-cut problem in a graph whose nodes are the matting components and whose edge weights represent matting penalty. Since the function we are trying to minimize is not submodular in general, we rewrite (14) as a sum of nonnegative local and pairwise terms  $J(\alpha) = E(b)$ , where  $b$  is the binary vector representing the components in  $\alpha$ . Alternatively, one could use algorithms for minimizing nonsubmodular functions such as the QPBO method (see [12] for a review on such methods).

We define the energy  $E(b)$  as

$$E(b) = \sum_k E_k(b^k) + \sum_{k,l} E_{k,l}(b^k - b^l)^2, \quad (15)$$




Fig. 6. A comparison of mattes produced by different matting methods from minimal user input (yellow scribbles indicate foreground, while green scribbles indicate background).

where  $E_k(0) = \infty$  if the  $k$ th component is constrained to belong to the foreground,  $E_k(1) = \infty$  if the component is constrained as background, and 0 otherwise. To define the pairwise term, we note that  $\phi^{k,k} = \sum_{l \neq k} \phi^{k,l}$  (this follows from the fact that  $L(\sum \alpha_k) = L\mathbf{1} = \mathbf{0}$  and, thus,  $\Phi\mathbf{1} = \mathbf{0}$ ) and, as a result,

$$b^T \Phi b = \sum_{k,l} -\phi^{k,l} (b^k - b^l)^2. \quad (16)$$

Therefore, if we define  $E_{k,l} = -\phi^{k,l}$ , we obtain that

$$\sum_{k,l} E_{k,l} (b^k - b^l)^2 = J(\alpha). \quad (17)$$

However, in order to search for a min-cut efficiently, we use a positive approximation and define  $E_{k,l} = \max(0, -\phi^{k,l})$ . To justify this approximation, we can prove that the approximation is exact in all local image windows for which no more than two components are active. When good components are extracted from an image, the majority of image windows will be associated with no more than two components (that is, no more than two components will be active). Indeed, we have empirically observed that, for most matting component pairs,  $\phi^{k,l} < 0$ .

**Claim 2.** *The correlation  $\phi^{k,l}$  between components  $\alpha_k, \alpha_l$  will be negative if every local image window  $w$  satisfies one of the following conditions:*

- c1.  $\alpha_i^k = 0 \forall i \in w$ .
- c2.  $\alpha_i^l = 0 \forall i \in w$ .
- c3.  $\alpha_i^k = 1 - \alpha_i^l \forall i \in w$ .

**Proof.** Following (6), we rewrite the correlation between components via the matting Laplacian as a sum of correlations over local windows:

$$\phi^{k,l} = \alpha^{kT} L \alpha^l = \sum_{q \in I} \alpha^{kT} L_q \alpha^l, \quad (18)$$

where  $L_q$  is an  $N \times N$  matrix if  $(i, j) \in w_q$   $L_q(i, j)$  is defined using (6) and  $L_q(i, j) = 0$  otherwise. We will show that, under the above conditions,  $\alpha^{kT} L_q \alpha^l < 0$  for

every image window  $w_q$ . This will follow immediately if  $\alpha_i^k = 0$  or  $\alpha_i^l = 0 \forall i \in w$ . For the third case, we note that

$$\alpha^{kT} L_q \alpha^l = \alpha^{kT} L_q (\mathbf{1} - \alpha^k) = -\alpha^{kT} L_q \alpha^k \leq^{***} 0, \quad (19)$$

where \* follows from the condition c3, \*\* follows from the fact that the matting cost of the constant vector is 0, and \*\*\* follows from the fact that each of the  $L_q$  matrices is positive semidefinite.  $\square$

Using the above graph formulation, finding the optimal foreground/background assignment does not involve an exponential search and is found efficiently, in time polynomial in the number of components, as a graph min-cut. As a result, if the matting components are precomputed, the optimal matte may be computed very rapidly, enabling interactive responses to user input. The computational challenges of our algorithm are equivalent to those of conventional spectral segmentation techniques. Specifically, it takes our unoptimized Matlab implementation a couple of minutes to compute the matting components for the images in Fig. 6. However, this preprocessing step can be done offline and, once the matting components are available, it only takes a few additional seconds to construct a matte given the user's constraints.

#### 4.2.2 Matte Extraction Using User Scribbles

Fig. 6 presents a few examples where a foreground matte was extracted from an image based on a small number of foreground (white) and background (black) markings provided by the user. The second column shows the resulting matte extracted by the approach described above (the scribbles are used to reduce the space of splitting hypotheses: a component is constrained to belonging to the foreground whenever its area contains a white scribble). The remaining columns show the mattes generated from the same input by a number of previous methods [14], [24], [8], [22]. None of these previous approaches is able to recover a reasonable matte from such minimal user input. In particular, although our approach uses the same matting Laplacian as [14], our results are very different from those obtained by directly minimizing the quadratic matting cost






Fig. 7. The middle region is not constrained and the method of Levin et al. assigns it an average nonopaque value.

(5) subject to user-specified constraints. The main drawback of such direct optimization is that, whenever an image contains distinct connected components without any constraints inside them, a quadratic cost such as (5) tends to assign them some average nonopaque values, as demonstrated by the simple example in Fig. 7. The core of this problem is that the quadratic cost of [14] places strong assumptions on the foreground and background distributions, but imposes no restrictions on  $\alpha$ . Thus, it searches for *continuous* solutions without taking into account that, for a mostly opaque foreground object, the matte should be strictly 0 or 1 over most of the image.

#### 4.2.3 Matte Extraction by Component Labeling

Once the matting components of an image have been computed, placing hard constraints by a set of scribbles or a trimap is not the only way for the user to specify her intent. The matting components suggest a new, more direct user interaction mode which was not possible until now: In this mode, the user is presented with the precomputed matting components and may simply label some of them as background or foreground. The labeled components then become constrained accordingly in the min-cut problem. The advantage of such an interface is illustrated in Fig. 8, where the large fuzzy hair areas do not lend themselves to placement of hard constraints. Thus, the best trimap we could practically expect leaves such areas unconstrained (Fig. 8c). The least-squares matte of [14] populates these areas with average gray values (Fig. 8d). In contrast, by searching for the cheapest assignment of matting components consistent with the trimap, we obtain the matte in Fig. 8e. In this case, no oversmoothing is observed, but some of the fuzzy hair was not selected to belong to the foreground. However, if the user is allowed to directly select three additional components (highlighted in red in Fig. 8g) as foreground, we obtain the matte in Fig. 8f.

## 5 QUANTITATIVE EVALUATION

To quantitatively evaluate our approach and compare it with previous methods, we captured ground truth data. Three different dolls were photographed in front of a computer monitor displaying seven different background images (Fig. 9a). A ground truth matte was then extracted for each doll using a least-squares framework [21]. Each image was downsampled to  $560 \times 820$  pixels and the tests described below were performed on (overlapping)  $200 \times 200$  windows cropped from these images. For our approach, 60 matting components were extracted using the 70 smallest eigenvectors of each cropped window. The running time of our

unoptimized Matlab implementation (on a 3.2 GHz CPU) was a few minutes for each  $200 \times 200$  window.

To design a comparison between matte extraction using matting components and previous matting algorithms, we need to address the two noncompatible interfaces and it is not clear how to measure the amount of user effort involved in each case. While previous approaches were designed to work with hard constraints (scribbles or trimap), our new approach enables a new interaction mode by component selection. Therefore, in our experiments, we attempted to determine how well each approach can do, given the best possible user input. Thus, we first used the ground truth matte to generate an “ideal” trimap. The unknown region in this trimap was constructed by taking all pixels whose ground truth matte values are between 0.05 and 0.95 and dilating the resulting region by four pixels. The resulting trimap was used as input for four previous matting algorithms: Levin et al. [14], Wang and Cohen [24], random walk matting [8], and Poisson matting [22]. We also ran our method twice in each experiment: 1) using the same trimap to provide a partial labeling of the matting components, followed by a min-cut computation, as described in Section 4.2 and 2) using the ground truth matte to select the subset of matting components that minimizes the distance of the resulting matte from the ground truth, thus simulating the ideal user input via the direct component picking interface. The SSD errors between the mattes produced by the different methods and the ground truth matte (averaged over the different backgrounds and the different windows) are plotted in Fig. 9b. It is apparent that, given a sufficiently precise trimap, our method offers no real advantage (when given the same trimap as input) over the least-squares matting of Levin et al., which produced the most numerically accurate mattes. However, when simulating the best labeling of components, our approach produced the most accurate mattes, on average.

While our experiment compares the quality of mattes produced from an ideal input, a more interesting comparison might be to measure the amount of user time required for extracting a satisfactory matte with each approach. Ideally, we would also like to measure whether (or to what degree) a component picking interface is more intuitive than a scribble-based interface. Such a comparison involves a nontrivial user study and is left for future work.

Given the strong analogy between spectral matting and hard spectral segmentation, we would like to gain some intuition about the possible advantage of using matting components versus standard hard segmentation components (also known as super-pixels). The answer, of course,




Fig. 8. Benefits of direct component labeling. (a) A test image. (b) Input. (c) Trimap. (d) Levin et al. [14] from trimap. (e) Components from trimap. (f) Component labeling. (g) Matting components.

depends on the application. If the final output is a hard segmentation, matting components probably do not offer an advantage over standard hard components. On the other hand, when the goal is a fuzzy matte, it is better to explicitly construct matting components, as we do, rather than first compute a hard segmentation and then feather the boundaries (as in [16], [19], for example). To show this, we compare the two approaches. We first extract matting

components from each of the test images and select the subset of matting components which will minimize the distance from the ground truth matte. The second approach is to select a subset of hard components (we used the available implementation of Yu and Shi [26]) that best approximates the ground truth. We then apply morphological operations (we have experimented with several constant radius erosion windows) on the resulting hard

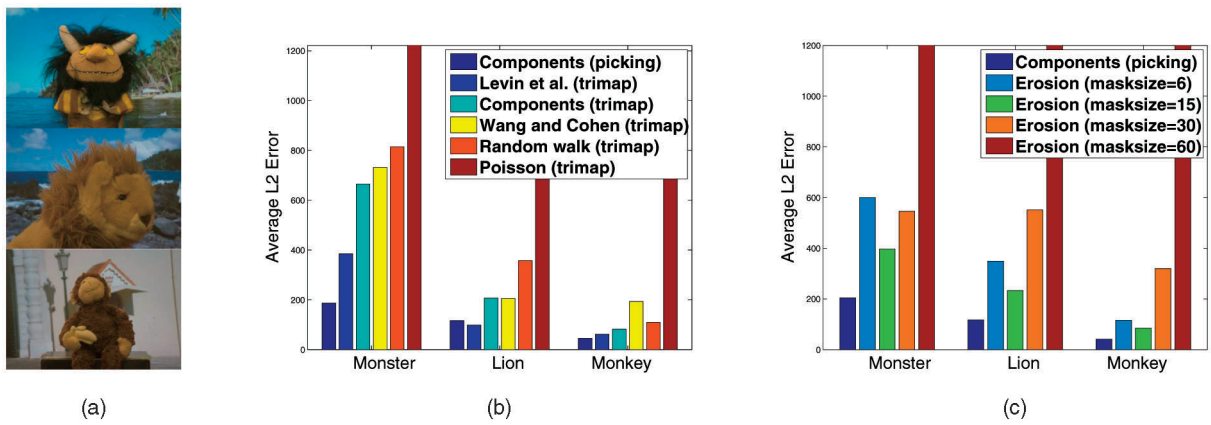


Fig. 9. Quantitative evaluation. (a) A few test images. (b) A comparison with other matting methods. (c) Spectral matting versus hard segmentation.

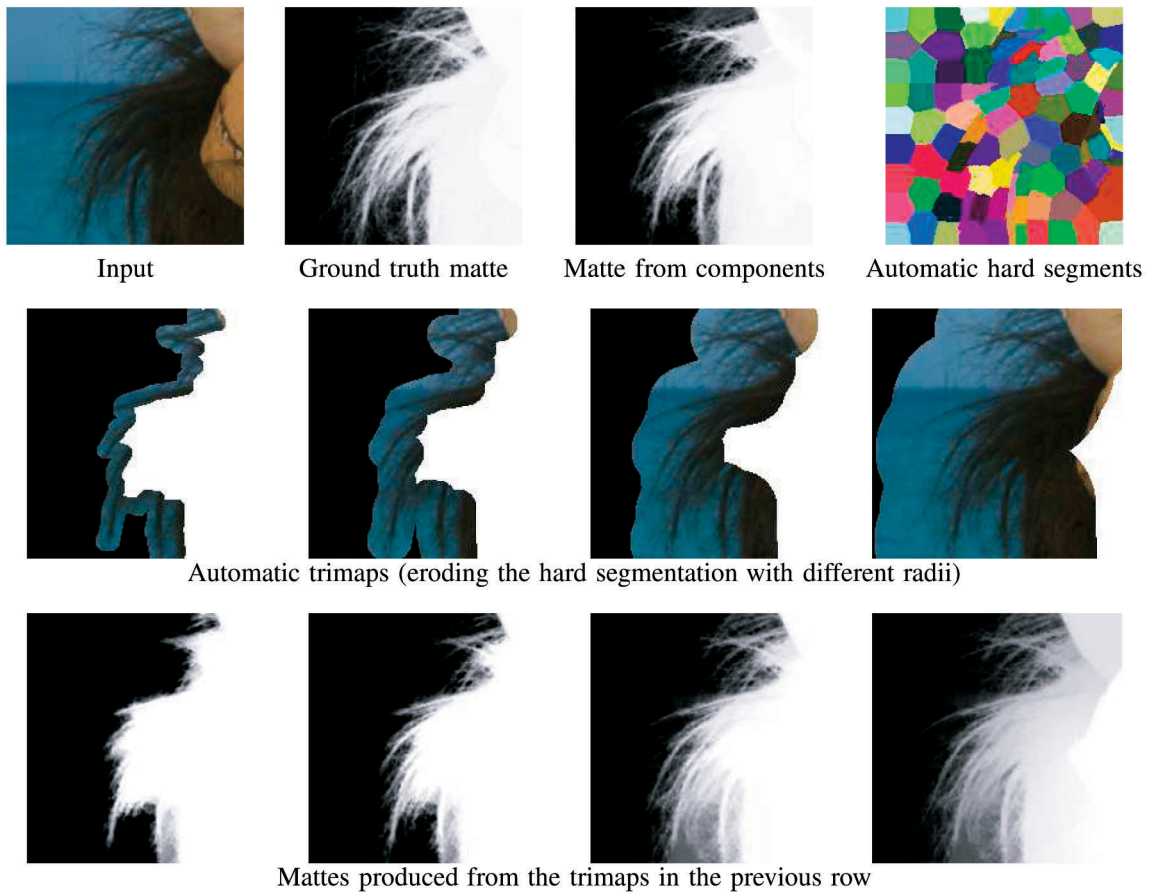


Fig. 10. Spectral matting versus obtaining trimaps from a hard segmentation.

mask, create a trimap, and run the matting algorithm of [14]. However, since the optimal radius of the erosion window strongly depends on the local image structure and varies over the image, it is impossible to obtain an ideal trimap with a constant radius window. This problem is illustrated visually in Fig. 10.

Fig. 9c shows the SSD errors (averaged over the different backgrounds and the different windows) of the two approaches, which indicate that optimally picking the matting components indeed results in more accurate mattes than those obtained by feathering a hard segmentation.

## 6 DISCUSSION

In this work, we have derived an analogy between standard (hard) spectral image segmentation and image matting and have shown how matting components may be automatically obtained from the smallest eigenvectors of the matting Laplacian. From a practical standpoint, matting components can help automate the matte extraction process and reduce user effort. Matting components also suggest a new mode of user control over the extracted matte: While, in previous methods, the result is controlled by placement of hard constraints in image areas where the matte is either completely opaque or completely transparent, our new




Fig. 11. Demonstration of limitations. Top: input image. Bottom: (a) Ground truth matte. Mattes from (b) 70 and (c) 400 eigenvectors.

approach may provide the user with a simple intuitive preview of optional outputs and thus enables the user to directly control the outcome in the fractional parts of the matte as well.

**Limitations.** Our method is most effective in automating the matte extraction process for images that consist of a modest number of visually distinct components. However, for highly cluttered images, component extraction proves to be a more challenging task. As an example, consider the case shown in Fig. 11. The input image consists of a large number of small components. Projecting the ground truth matte (Fig. 11a) on the subspace spanned by the 70 smallest eigenvectors results in a poor approximation (Fig. 11b). Recall that, since the matting components are obtained via a linear combination of the eigenvectors, they can do no better than the eigenvectors themselves and, thus, Fig. 11b is the best matte that we could hope to construct from up to 70 matting components. Thus, it is quite clear that this number of components is insufficient to produce an accurate matte for this image. A better matte may be obtained from the 400 smallest eigenvectors (Fig. 11c), but even this matte leaves room for improvement. We have not been able to test more than 400 eigenvectors due to computational limitations. We have empirically observed that this problem is significantly reduced if matting components are computed in local image windows independently and we are currently investigating methods for stitching together components obtained in different windows.

One major challenge in spectral matting is determining the appropriate number of matting components for a given image. This is a fundamental difficulty shared by all spectral segmentation methods. While the question of automatically selecting the number of components has been investigated (e.g., [27]), this parameter is still often manually adjusted. For the applications described in this

paper, we found that a useful strategy is to oversegment the image and group the components later using additional cues. A second free parameter in the algorithm is the number of smallest eigenvectors from which the components are formed (the number should be larger than or equal to the number of components). In practice, we have observed that the performance is not very sensitive to this number and all results in this paper were obtained using the 70 smallest eigenvectors. Fig. 12 demonstrates the effect of varying the number of matting components and the number of eigenvectors. It can be seen that usually it is better to use a large number of matting components and a large number of eigenvectors, but this obviously makes the algorithm slower. When too many matting components are used, some regions in the alpha matte may become too transparent. This is due to some matting components that were not grouped correctly by the min-cut algorithm. The algorithm is much less sensitive to the number of eigenvectors being used. Using 50 to 200 eigenvectors usually produced similar results. It should be noted that increasing the number of eigenvectors makes the resulting alpha matte more opaque as it gives more freedom to the nonlinear optimization, which looks for opaque solutions. (In the extreme case, using all eigenvectors will produce binary mattes.)

**Future directions.** An important potential advantage of presegmenting the image into matting components is the option to compute meaningful color or texture histograms, or other statistics, within each component. The histogram similarity can provide another important cue to guide component grouping. This ability might significantly improve matting algorithms that make use of color models, such as [24]. For example, the current strategy in [24] is to build an initial color model using only the small number of pixels under the scribbles. This simple initialization is




Fig. 12. The effect of changing the number of eigenvectors and the number of matting components on the resulting matte.

known to make the algorithm sensitive to small shifts in the scribble location.

Given the growing interest in the matting problem and the large amount of recent matting research, it seems that an important future challenge is the design of an appropriate comparison between different user interaction modes and different matting algorithms. The ground truth data collected in this work is a step toward this goal, yet a proper user study is required in order to evaluate the amount of user time required for producing good results with each method.

Our code and ground truth data are available at: [www.vision.huji.ac.il/SpectralMatting](http://www.vision.huji.ac.il/SpectralMatting).

## REFERENCES

- [1] E. Borenstein and S. Ullman, "Class-Specific, Top-Down Segmentation," *Proc. Seventh European Conf. Computer Vision*, pp. 109-122, 2002.
- [2] Y. Chuang, B. Curless, D. Salesin, and R. Szeliski, "A Bayesian Approach to Digital Matting," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pp. 264-271, 2001.
- [3] J. Dennis, J. Robert, and B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Prentice-Hall, 1983.
- [4] W. Donath and A. Hoffman, *Algorithms for Partitioning Graphs and Computer Logic Based on Eigenvectors of Connection Matrices*, 1973.
- [5] M. Fiedler, "Algebraic Connectivity of Graphs," *Czechoslovak Math. J.*, vol. 23, no. 98, pp. 298-305, 1973.
- [6] M. Fiedler, "Eigenvectors of Acyclic Matrices," *Czechoslovak Math. J.*, vol. 25, no. 100, pp. 607-618, 1975.
- [7] M. Fiedler, "A Property of Eigenvectors of Nonnegative Symmetric Matrices and Its Application to Graph Theory," *Czechoslovak Math. J.*, vol. 25, no. 100, pp. 619-633, 1975.
- [8] L. Grady, T. Schiwietz, S. Aharon, and R. Westermann, "Random Walks for Interactive Alpha-Matting," *Proc. Fifth IASTED Int'l Conf. Visualization, Imaging and Image Processing*, pp. 423-429, 2005.
- [9] Y. Guan, W. Chen, X. Liang, Z. Ding, and Q. Peng, "Easy Matting: A Stroke Based Approach for Continuous Image Matting," *Computer Graphics Forum, Proc. Eurographics '06*, vol. 25, no. 3, pp. 567-576, 2006.
- [10] S. Guattery and G.L. Miller, "On the Quality of Spectral Separators," *SIAM J. Matrix Analysis and Applications*, vol. 19, no. 3, pp. 701-719, 1998.
- [11] K. Hall, "An R-Dimensional Quadratic Placement Algorithm," *Management Science*, vol. 17, no. 3, pp. 219-229, 1970.
- [12] V. Kolmogorov and C. Rother, "Minimizing Nonsubmodular Functions with Graph Cuts—A Review," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 29, no. 7, pp. 1274-1279, July 2007.
- [13] K. Lang, "Fixing Two Weaknesses of the Spectral Method," *Proc. Advances in Neural Information Processing Systems*, vol. 18, 2005.
- [14] A. Levin, D. Lischinski, and Y. Weiss, "A Closed Form Solution to Natural Image Matting," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pp. 61-68, 2006.
- [15] A. Levin, A. Rav-Acha, and D. Lischinski, "Spectral Matting," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, June 2007.
- [16] Y. Li, J. Sun, C.-K. Tang, and H.-Y. Shum, "Lazy Snapping," *ACM Trans. Graphics*, vol. 23, no. 3, pp. 303-308, 2004.
- [17] A. Ng, M. Jordan, and Y. Weiss, "On Spectral Clustering: Analysis and an Algorithm," *Proc. Advances in Neural Information Processing Systems*, 2001.
- [18] A. Pothen, H. Simon, and K.-P. Liou, "Partitioning Sparse Matrices with Eigenvalues of Graphs," *SIAM J. Matrix Analysis and Applications*, vol. 11, no. 3, pp. 430-452, 1990.
- [19] C. Rother, V. Kolmogorov, and A. Blake, "'GrabCut': Interactive Foreground Extraction Using Iterated Graph Cuts," *ACM Trans. Graphics*, vol. 23, no. 3, pp. 309-314, 2004.
- [20] J. Shi and J. Malik, "Normalized Cuts and Image Segmentation," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888-905, Aug. 2000.

- [21] A. Smith and J. Blinn, "Blue Screen Matting," *Proc. ACM SIGGRAPH '96*, pp. 259-268, 1996.
- [22] J. Sun, J. Jia, C.-K. Tang, and H.-Y. Shum, "Poisson Matting," *ACM Trans. Graphics*, vol. 23, no. 3, pp. 315-321, 2004.
- [23] D. Tolliver and G. Miller, "Graph Partitioning by Spectral Rounding: Applications in Image Segmentation and Clustering," *Proc. IEEE Int'l Conf. Computer Vision and Pattern Recognition*, pp. 1053-1060, 2006.
- [24] J. Wang and M. Cohen, "An Iterative Optimization Approach for Unified Image Segmentation and Matting," *Proc. Int'l Conf. Computer Vision*, pp. 936-943, 2005.
- [25] Y. Weiss, "Segmentation Using Eigenvectors: A Unifying View," *Proc. Int'l Conf. Computer Vision*, pp. 975-982, 1999.
- [26] S.X. Yu and J. Shi, "Multiclass Spectral Clustering," *Proc. Int'l Conf. Computer Vision*, pp. 313-319, 2003.
- [27] L. Zelnik-Manor and P. Perona, "Self-Tuning Spectral Clustering," *Proc. Advances in Neural Information Processing Systems*, 2005.



**Anat Levin** received the BSc, MSc, and PhD degrees in computer science from the Hebrew University of Jerusalem. She is a postdoctoral associate at the Massachusetts Institute of Technology's Computer Science and Artificial Intelligence Laboratory. In 2009, she will join the Weizmann Institute of Science's computer science faculty. Her research interests include computer vision, computer graphics, and machine learning. In particular, she has worked on

low and mid-level vision, computational photography, and visual recognition.



**Alex Rav-Acha** received the BSc, MSc, and PhD degrees in computer science from the Hebrew University of Jerusalem. He is currently a postdoctoral associate at the Weizmann Institute of Science. His research interests include video summary, video editing, image deblurring, and motion analysis.



**Dani Lischinski** received the BSc and MSc degrees in computer science from the Hebrew University of Jerusalem in 1987 and 1989, respectively, and the PhD degree in computer science from Cornell University in 1994. He is currently on the faculty of the School of Engineering and Computer Science at the Hebrew University of Jerusalem. His research interests include computer graphics and image and video processing. In particular, he has

worked on algorithms for photorealistic image synthesis, image-based modeling and rendering, texture synthesis, and editing, tone mapping, and manipulation of images and video.

▷ **For more information on this or any other computing topic, please visit our Digital Library at [www.computer.org/publications/dlib](http://www.computer.org/publications/dlib).**