

Spectral Matting

Anat Levin^{1,2}

Alex Rav-Acha¹

Dani Lischinski¹

¹School of CS & Eng
The Hebrew University

²CSAIL
MIT

Abstract

We present spectral matting: a new approach to natural image matting that automatically computes a set of fundamental 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.

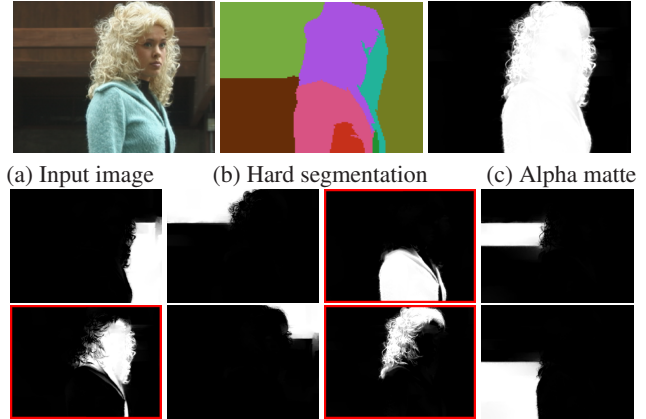
1. Introduction

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 under-constrained, all of the existing methods require the user to provide additional constraints in the form of a trimap [2, 14, 4] or a set of brush strokes [16, 7, 5]. Thus, the question of whether (or to what degree) is it 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 [12, 17, 10, 18]. These methods perform unsupervised image segmentation by examining the smallest eigenvectors of the image's graph Laplacian matrix. This work, for the first time, extends this idea from producing hard segments to soft *matting components*.

Spectral segmentation methods, such as [12], 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 attempt 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.* [7]. Once obtained, these matting components serve as building blocks for construction of complete foreground mattes.



(d) Matting components computed by our method.

Figure 1. Spectral segmentation and spectral matting

This concept is illustrated in Figure 1. Given the input image in Figure 1a, one can produce an unsupervised disjoint hard partitioning of the image using, e.g., [18] (Figure 1b). In contrast, we compute a set of overlapping, fractional, matting components, visualized in Figure 1d. Combining three of these components (framed in red) yields the foreground matte of the girl, shown in Figure 1c.

In summary, our main contribution is the introduction of the concept of fundamental matting components and the resulting first unsupervised matting algorithm. Of course, just like unsupervised segmentation, unsupervised matting is an ill-posed problem. Thus, we also describe two extensions that use these fundamental matting components to construct a particular matte: (i) present the user with several matting alternatives to choose from; or (ii) let the user specify her intent by just a few mouse clicks.

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 α^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 non-negative and sum to 1 at every pixel. The intuitive motivation for having these components is that similarly to the individual low-level fragments in an over-segmented image they may be used to construct higher level, semantically meaningful foreground mattes, as demonstrated in Figure 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 [7]. 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 [12, 17, 10, 18]. 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.¹

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 an eigenvector of L with eigenvalue 0.

Now suppose that the image consists of K connected components C^1, \dots, C^K such that $\{1, \dots, N\} = \bigcup_{k=1}^K C^k$,

¹In 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.

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 eigenvectors of L with eigenvalue 0. However, computing the 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 eigenvectors with zero eigenvalue. 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 [10, 18, 15, 19, 6]. The simplest approach [10] 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 [18, 19], 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.

The matting Laplacian was introduced by Levin *et al.* [7] in order to evaluate the quality of a matte without explicitly estimating the foreground and background colors in eq. (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 α 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 :

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

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

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

Here 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. $L(i, j)$ is defined as:

$$\sum_{q|(i,j) \in w_q} \left(\delta_{ij} - \frac{1}{|w_q|} \left(1 + (I_i - \mu_q)^T (\Sigma_q + \frac{\varepsilon}{|w_q|} I_3)^{-1} (I_j - \mu_q) \right) \right) \quad (6)$$

Here δ_{ij} is the Kronecker delta, μ_q is the 3×1 mean color vector in the window w_q around pixel q , Σ_q is a 3×3 covariance matrix in the same window, $|w_q|$ is the number of pixels in the window, and I_3 is the 3×3 identity matrix.

The cost (5) has a trivial minimum which is a constant α vector, and thus in the user assisted framework described in [7], $J(\alpha)$ is minimized subject to user constraints. Levin *et al.* observe that the smallest eigenvectors of the matting Laplacian (6) capture information about the fuzzy cluster assignments of pixels in the image, even before any user-specified constraints are taken into account. However, they make 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 span the individual matting components of the image.

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 under reasonable conditions, the actual matting components belong to the nullspace of the matting Laplacian. We say that a matting component α^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 color of each layer must be 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 eq. 6 with $\varepsilon = 0$) if every local image window w satisfies one of the following conditions:*

1. A single component α^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 $\alpha^{k_1}, \alpha^{k_2}, \alpha^{k_3}$ 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 (eq. 4). Thus, in order to show that a matte component α^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



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

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 one at every image pixel, the single active component α^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 theorem 2 in [7].

Case 3: Since $F^{k_1}, F^{k_2}, F^{k_3}$ are constant within w and their colors are linearly independent, there exist a^R, a^G, a^B and b such that $\langle a, F^{k_1} \rangle + b = 1, \langle a, F^{k_2} \rangle + b = 0$ and $\langle a, F^{k_3} \rangle + b = 0$. As $I = \alpha^{k_1} F^{k_1} + \alpha^{k_2} F^{k_2} + \alpha^{k_3} F^{k_3}$ we get that $\langle a, I \rangle + b = \alpha^{k_1}$, so that α^{k_1} is a linear function of the image. A similar argument holds for α^{k_2} and α^{k_3} . \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 zero eigenvectors of L .

In most real images, the assumptions of claim 1 don't hold exactly, and thus the matting Laplacian might not have multiple eigenvectors whose eigenvalue is 0. Yet if the layers are sufficiently distinct, they are generally captured by the smallest eigenvectors of L . For example, Figure 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 Figure 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. 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

$$\sum_{i,k} |\alpha_i^k|^\gamma + |1 - \alpha_i^k|^\gamma, \text{ where } \alpha^k = Ey^k \quad (7)$$

subject to $\sum_k \alpha_i^k = 1$

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. Without the requirement $\alpha^k = Ey^k$ the sparsity term would be minimized by binary vectors, but as the vectors α^k are restricted to linear combinations of the eigenvectors they must maintain the fuzzy layer boundaries. Although we do not explicitly constrain the α 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 non-convex one and we optimize it iteratively using Newton's method [3] by constructing a sequence of second order approximations (whose minimization involves the solution of a $K^2 \times K^2$ linear system). More details may be found in [8].

Since the cost (7) is not convex, the result of the Newton process strongly depends on the quality of 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 :

$$\alpha^k = EE^T m^{C^k}. \quad (8)$$

It can be shown that the resulting matting components sum to one and thus provide a legal solution for eq. (7).

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).

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. Suppose $\alpha^{k_1}, \dots, \alpha^{k_n}$ were designated as foreground components, then the complete foreground matte is obtained simply by adding them together:

$$\alpha = \alpha^{k_1} + \dots + \alpha^{k_n} \quad (9)$$

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

For the applications discussed below, one would like to compare multiple grouping hypotheses, and thus measure the quality of the resulting α -matte as $J(\alpha) = \alpha^T L \alpha$, 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 pre-compute the correlations between the matting components via L and store them in a $K \times K$ matrix Φ , where

$$\Phi(k, l) = \alpha^k{}^T L \alpha^l. \quad (10)$$

The matte cost can then be computed as

$$J(\alpha) = b^T \Phi b, \quad (11)$$

where b is a K dimensional binary vector indicating the selected components. Thus, if Φ has been pre-computed, $J(\alpha)$ 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 by eq. (11). However, the matting cost is usually biased toward mattes which assign non constant 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 suggest several criteria which overcome this bias. One approach is to search for quotient cuts (*e.g.*, normalized cuts [12]) 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 [6] 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 which assign at least 30% of the pixels to the foreground and at least 30% 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 eq. (11).

Figure 3 shows some results produced by the unsupervised matting approach described above. In each of these examples the hypothesis with the highest score indeed corresponds to the "correct" foreground matte, but some of the other hypotheses are quite sensible as well, considering that our approach does not attempt to perform any high-level image understanding. Of course, it isn't hard to find examples where unsupervised matting fails. For example, whenever the foreground or background objects consist of several visually distinct components, the assignment with the minimal matting cost might not correspond to our vi-




Figure 3. Unsupervised matting results for two images. The hypotheses are ordered according to their score.

sual perception. In fact, it is well known within the image segmentation community that while unsupervised bottom-up cues can 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 [11], or motion and focus cues. In the remainder of this paper, however, we focus on user-guided matting instead.

4.2. User-Guided Matting

Matting components can also be quite useful in 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 approximate the matting cost (11) 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 (see [8] for details). In this formulation, finding the optimal assignment does not involve an exponential search and is found efficiently in time polynomial in the number of components. As a result, if the matting components are pre-computed, 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 Figure 4. However, this pre-processing step can be done offline, and once the matting components are available, it only takes an additional few seconds to con-




Figure 5. The middle region is not constrained, and the method of Levin *et al.* assigns it an average non-opaque value.

struct a matte given the user’s constraints.

Figure 4 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 belong 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 [7, 16, 4, 14]. 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 [7], our results are very different from those obtained by directly minimizing the quadratic matting cost (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 non-opaque values, as demonstrated by the simple example in Figure 5. The core of this problem is that the quadratic cost of [7] places strong assumptions on the foreground and background distributions, but imposes no restrictions on α . 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.

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 wasn’t possible until now: in this




Figure 4. A comparison of mattes produced by different matting methods from minimal user input.

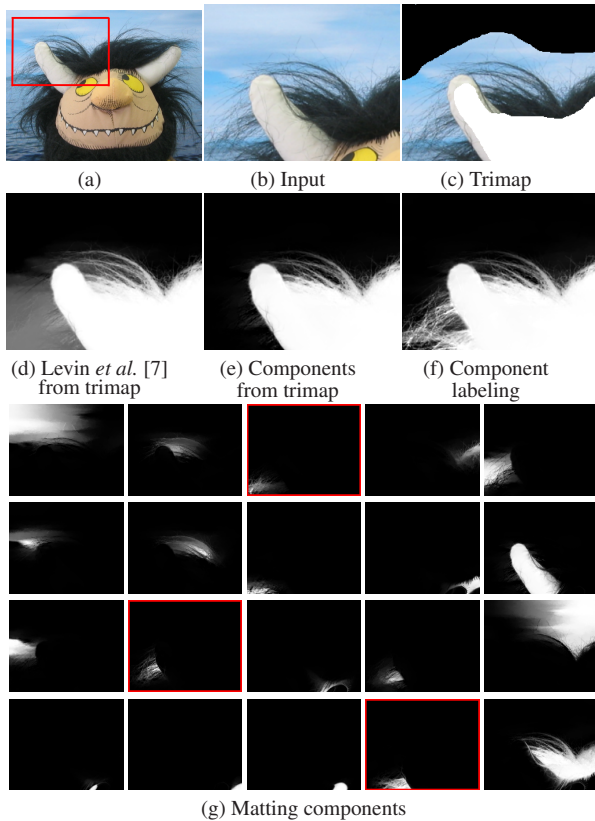


Figure 6. Benefits of direct component labeling.

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 Figure 6, 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

(Figure 6c). The least squares matte of [7] populates these areas with average gray values (Figure 6d). In contrast, by searching for the cheapest assignment of matting components consistent with the trimap, we obtain the matte in Figure 6e. In this case no over-smoothing 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 Figure 6g) as foreground, we obtain the matte in Figure 6f.

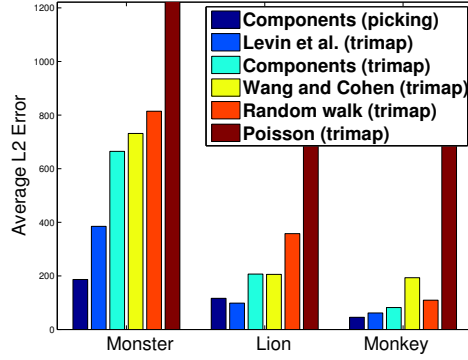
5. Quantitative evaluation

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 (Figure 7a). A ground truth matte was then extracted for each doll using a least squares framework [13]. Each image was downsampled to 560×820 pixels, and the tests described below were performed on (overlapping) 200×200 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 3.2GHz CPU) was a few minutes for each 200×200 window.

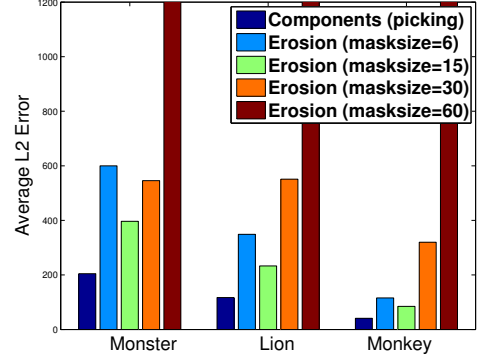
To design a comparison between matte extraction using matting components and previous matting algorithms we need to address the two non compatible interfaces, and it is not clear how to measure the amount of user effort involved in each case. While previous approaches were designed to work with hard 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 can each approach 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



(a) Three of the test images



(b) A comparison with other matting methods



(c) Spectral matting vs. hard segmentation

Figure 7. Quantitative evaluation

the resulting region by 4 pixels. The resulting trimap was used as input for four previous matting algorithms: Levin *et al.* [7], Wang and Cohen [16], random walk matting [4], and Poisson matting [14]. We also ran our method twice in each experiment: (i) 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 (ii) 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 Figure 7b. 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 non trivial 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, 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 [11, 9], 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 [18]) that best approximates the ground truth. We then apply morphological operations (we have experimented with several constant radius erosion windows) on the resulting hard mask, create a trimap and run the matting algorithm of [7]. 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 the supplementary materials. Figure 7c 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 hard spectral image segmentation and image matting, and have shown how fundamental matting components may be automatically obtained from the smallest eigenvectors of the matting Laplacian. This is a new interesting theoretical result, establishing a link between two previously independent research areas. From the 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 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,




Figure 8. Limitations. Top: input image; Bottom: Ground truth matte (a); Mattes from 70 (b) and 400 (c) eigenvectors.

for highly cluttered images, component extraction proves to be a more challenging task. For example, consider the example in Figure 8. The input image consists of a large number of small components. Projecting the ground truth matte (Figure 8a) on the subspace spanned by the 70 smallest eigenvectors results in a poor approximation (Figure 8b). 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 Figure 8b 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 (Figure 8c), 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. 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 component has been investigated (e.g. [19]), this parameter is still often manually adjusted. For the applications described in this paper we found that a useful strategy is to over-segment 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 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.

Future directions: An important potential advantage of pre-segmenting 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 which make use of color models such

as [16]. For example, the current strategy in [16] is to build an initial color model using only the small number of pixels under the scribbles. This poor initialization is 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

References

- [1] E. Borenstein and S. Ullman. Class-specific, top-down segmentation. In *ECCV*, 2002.
- [2] Y. Chuang, B. Curless, D. Salesin, and R. Szeliski. A Bayesian approach to digital matting. In *CVPR*, 2001.
- [3] J. Dennis, J. Robert, and B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Prentice-Hall, 1983.
- [4] L. Grady, T. Schiwietz, S. Aharon, and R. Westermann. Random walks for interactive alpha-matting. In *VIIP*, 2005.
- [5] Y. Guan, W. Chen, X. Liang, Z. Ding, and Q. Peng. Easy matting. In *Eurographics*, 2006.
- [6] K. Lang. Fixing two weaknesses of the spectral method. In *NIPS*. 2005.
- [7] A. Levin, D. Lischinski, and Y. Weiss. A closed form solution to natural image matting. In *CVPR*, 2006.
- [8] A. Levin, A. Rav-Acha, and D. Lischinski. Spectral matting. Hebrew University Technical Report, May 2007.
- [9] Y. Li, J. Sun, C.-K. Tang, and H.-Y. Shum. Lazy snapping. *ACM Trans. Graph.*, 23(3):303–308, 2004.
- [10] A. Ng, M. Jordan, and Y. Weiss. On spectral clustering: Analysis and an algorithm. In *NIPS*. 2001.
- [11] C. Rother, V. Kolmogorov, and A. Blake. “GrabCut”: interactive foreground extraction using iterated graph cuts. *ACM Trans. Graph.*, 23(3):309–314, 2004.
- [12] J. Shi and J. Malik. Normalized cuts and image segmentation. *PAMI*, 22(8):888–905, 2000.
- [13] A. Smith and J. Blinn. Blue screen matting. In *Proceedings of ACM Siggraph*, pages 259–268, 1996.
- [14] J. Sun, J. Jia, C.-K. Tang, and H.-Y. Shum. Poisson matting. *ACM Trans. Graph.*, 23(3):315–321, 2004.
- [15] D. Tolliver and G. Miller. Graph partitioning by spectral rounding: Applications in image segmentation and clustering. In *CVPR*, 2006.
- [16] J. Wang and M. Cohen. An iterative optimization approach for unified image segmentation and matting. In *ICCV*, 2005.
- [17] Y. Weiss. Segmentation using eigenvectors: A unifying view. In *ICCV*, pages 975–982, 1999.
- [18] S. X. Yu and J. Shi. Multiclass spectral clustering. In *ICCV*, pages 313–319, 2003.
- [19] L. Zelnik-Manor and P. Perona. Self-tuning spectral clustering. In *NIPS*. 2005.