Poisson Matting

Jian Sun[1] Jiaya Jia[2*] Chi-Keung Tang[2] Heung-Yeung Shum[1]

[1]Microsoft Research Asia [2]Hong Kong University of Science and Technology



Figure 1: Pulling of matte from a complex scene. From left to right: a complex natural image for existing matting techniques where the color background is complex, a high quality matte generated by Poisson matting, a composite image with the extracted koala and a constant-color background, and a composite image with the extracted koala and a different background.

Abstract

In this paper, we formulate the problem of natural image matting as one of solving Poisson equations with the matte gradient field. Our approach, which we call *Poisson matting*, has the following advantages. First, the matte is directly reconstructed from a continuous matte gradient field by solving Poisson equations using boundary information from a user-supplied trimap. Second, by interactively manipulating the matte gradient field using a number of filtering tools, the user can further improve Poisson matting results locally until he or she is satisfied. The modified local result is seamlessly integrated into the final result. Experiments on many complex natural images demonstrate that Poisson matting can generate good matting results that are not possible using existing matting techniques.

Keywords: matting, alpha channel, image compositing, Poisson equation.

1 Introduction

In image composition, a new image $I(x,y)$ can be blended from a background image $B(x,y)$ and a foreground image $F(x,y)$ with its alpha matte $\alpha(x,y)$ by the matting equation ((x,y) arguments deleted for clarity):

$$I = \alpha F + (1-\alpha)B \tag{1}$$

On the other hand, separation of α, F and B from a given image I is called the *pulling of matte* problem, or simply matting. In blue screen matting, α and F need to be reconstructed but B is known from a user-controlled environment. In natural image matting, all variables α, F and B need to be estimated.

Matting is inherently under-constrained because the matting equation has too many unknowns. Therefore, user interaction is essential to obtain good mattes. For example, in natural image matting (e.g., Bayesian matting [Chuang et al. 2001]), the user is required to supply a *trimap* that partitions the image into three regions: "definitely foreground", "definitely background" and "unknown region". In the unknown region, the matte can be estimated using the color statistics in the known foreground and background regions.

Indeed, as mentioned in [Smith and Blinn 1996], the key ideas to the success of blue-screen matting in many applications are the inclusion of a human in the environment setting cycle to determine whether a result is "a correct matte when s/he sees one", and providing a sufficiently rich set of controls for adjusting the result.

Adjusting the matting result in natural image matting is however not a straightforward operation. Most natural image matting approaches to date [Berman et al. 2000; Ruzon and Tomasi 2000; Hillman et al. 2001; Chuang et al. 2001] rely on sampling pixels in the known background and foreground. These samples are then used to estimate the matte in an unknown region in a statistically meaningful way. Results from these approaches may not be improved any further once the trimap has been carefully specified. For example, in Figure 1, the foreground hairs and the background branches are easily confused, especially in regions of low contrast. Without additional information, it is very difficult to produce a good matting result. Editing mattes directly at the pixel level would be tedious and impractical.

In this paper, we propose *Poisson matting* for natural images of complex scenes. Unlike previous methods which optimize a pixel's alpha, background and foreground colors in a statistical manner, our method operates directly on the gradient of the matte. This reduces the error caused by mis-classification of color samples in a complex scene. Poisson matting estimates the gradient of matte from the image, then reconstructs the matte by solving Poisson equations.

Our formulation is based on the assumption that intensity change in the foreground and background is smooth. We propose *global Poisson matting*, a semi-automatic approach to approximate matte from an image gradient given a user-supplied trimap. Foreground

*This research was done when Jiaya Jia was visiting MSR Asia.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2004 ACM 0730-0301/04/0800-0315 $5.00

Google
Exhibit 1013

and background colors can then be more robustly estimated.

More importantly, when global Poisson matting fails to produce high quality mattes due to a complex background, we introduce *local Poisson matting*, which manipulates a continuous gradient field in a local region. Local Poisson matting brings human interaction into the "pulling the matte" loop. In most cases, image gradients caused by foreground and background colors are visually distinguishable locally. Knowledge from the user can be effectively integrated into Poisson matting by using a set of tools that operates on the gradient field of the matte. As a result, our method can maintain the continuities of thin long threadlike shapes of foreground objects in a complex scene as shown in Figure 1. Moreover, after local operations, user modifications in a gradient field can be seamlessly propagated into the matte.

2 Related Work

Natural image matting. Several natural image matting methods [Berman et al. 2000; Ruzon and Tomasi 2000; Hillman et al. 2001; Chuang et al. 2001] have been proposed recently. Most of them contain two processes: 1) sample gathering: statistical information or samples of color F and B for each pixel in "unknown region" are collected from "definitely foreground" and "definitely background". 2) matte estimation: the matte is estimated for each pixel, given values of F and B by solving the matting equation.

In Knockout [Berman et al. 2000], the estimated F and B are weighted averages of the pixels along the perimeter of the known foreground and background regions. The final estimated α is also a weighted average of the matte intensity calculated in the RGB channels. In [Ruzon and Tomasi 2000] and [Hillman et al. 2001], the color samples for F and B are analyzed by a mixture of un-oriented Gaussians and principal components analysis (PCA) respectively. Bayesian matting [Chuang et al. 2001] first clusters the color samples for F and B. Each cluster is fitted with an oriented Gaussian distribution. A *maximum a posterior* (MAP) estimation of α, F and B is calculated simultaneously for each of the foreground and background pair in a Bayesian framework. The final α is chosen from the pair that produces the maximum likelihood. To date, Bayesian matting produces the best results in many cases. A detailed comparison and survey on natural image matting algorithms can be found in [Chuang et al. 2001].

However, these methods rely on color sampling which can be error-prone in complex scenes. Sampling wrong colors inevitably leads to poor matting results. Moreover, when the trimap can no longer be refined, it is unclear how these color sampling methods can continue to improve matting results, either automatically or interactively.

Other matting techniques. In [Smith and Blinn 1996], a triangulation solution is proposed to convert matting into an overconstrained problem. But it requires two shots of the foreground image and two different known backgrounds (colors). Difference matting [Qian and Sezan 1999] also needs two images: one with and the other without foreground. The difference of two images is mapped to an matte. Video matting [Chuang et al. 2002] is an extension of Bayesian matting. A bi-directional optical flow algorithm is used to interpolate the trimap in a video sequence, from user supplied key frames.

The Poisson equations has been used previously in tone mapping [Fattal et al. 2002], shadow removal [Finlayson et al. 2002] and image editing [Elder and Goldberg 2001; Pérez et al. 2003]. An image can be modified by manipulating the gradient field of an image automatically or interactively and then solving Poisson equations. We apply Poisson equations to image matting in this paper.

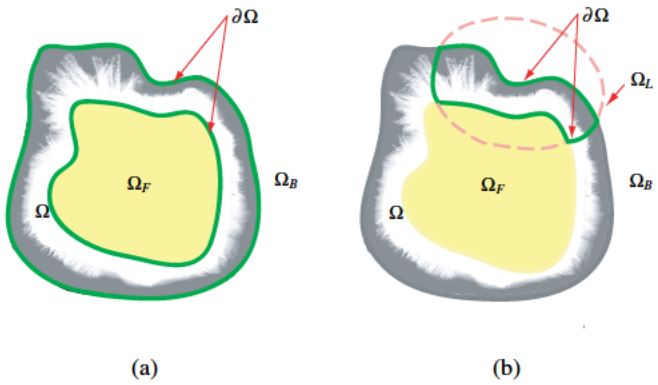


Figure 2: Boundary condition for Poisson matting. (a) Global Poisson matting: the trimap $\{\Omega_F, \Omega_B, \Omega\}$ is specified by the user. $\partial\Omega$ is the exterior boundary of unknown region Ω. (b) Local Poisson matting: the user selects a local region Ω_L interactively. $\partial\Omega$ is the exterior boundary of local unknown region $\Omega \cap \Omega_L$.

3 Poisson Matting

Poisson matting consists of two steps. First, an approximate gradient field of matte is computed from the input image. Second, the matte is obtained from its gradient field by solving Poisson equations.

In order to get an approximate gradient field of matte, we take the partial derivatives on both sides of the matting equation:

$$\nabla I = (F-B)\nabla\alpha + \alpha\nabla F + (1-\alpha)\nabla B \tag{2}$$

where $\nabla = (\frac{\partial}{\partial x}, \frac{\partial}{\partial y})$ is the gradient operator. This is the differential form of the matting equation, for R, G and B channels individually. In situations in which foreground F and background B are smooth, i.e., $\alpha\nabla F + (1-\alpha)\nabla B$ is relatively small with respect to $(F-B)\nabla\alpha$, we can get an approximate matte gradient field as follows:

$$\nabla\alpha \approx \frac{1}{F-B}\nabla I \tag{3}$$

It means that the matte gradient is proportional to the image gradient. This approximation of the gradient of the matte first appeared in [Mitsunaga et al. 1995]. To estimate the opacity around the boundaries of a solid object, they integrated the gradient of matte along a 1D path that is perpendicular to the boundary of the object. Using the same approximation, however, we reconstruct the matte more efficiently by solving Poisson equations in a 2D image space directly.

3.1 Global Poission matting

As shown in Figure 2(a), Ω_F, Ω_B and Ω are defined as "definitely foreground", "definitely background" and "unknown" regions respectively. For each pixel $p = (x, y)$ in the image, I_p is its intensity, F_p and B_p are the foreground and background intensity respectively. Let N_p be the set of its 4 neighbors. $\partial\Omega = \{p \in \Omega_F \cup \Omega_B | N_p \cap \Omega \neq \emptyset\}$ is the exterior boundary of Ω.

To recover the matte in the unknown region Ω given an approximate $(F-B)$ and image gradient ∇I, we minimize the following variational problem:

$$\alpha^* = \arg\min_{\alpha} \int\int_{p\in\Omega} ||\nabla\alpha_p - \frac{1}{F_p - B_p}\nabla I_p||^2 dp \tag{4}$$

with Dirichlet boundary condition $\alpha|_{\partial\Omega} = \widehat{\alpha}|_{\partial\Omega}$. We define

$$\widehat{\alpha}_p|_{\partial\Omega} = \begin{cases} 1 & p \in \Omega_F \\ 0 & p \in \Omega_B \end{cases} \quad (5)$$

This definition is consistent with a user-supplied trimap. The associated Poisson equations with the same boundary condition is:

$$\Delta\alpha = div(\frac{\nabla I}{F-B}) \quad (6)$$

where $\Delta = (\frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2})$ and div are Laplacian and Divergence operators respectively. Obtaining the unique solution of Poisson equations is a well studied problem. We use the Gauss-Seidel iteration with overrelaxation method (also used in [Pérez et al. 2003]). For color images, both $(F-B)$ and ∇I are measured in the grayscale channel.

Iterative optimization. Global Poisson matting is an iterative optimization process:

1. $(F-B)$ **initialization** Absolute values of F and B are not necessary, since $(F-B)$ provides enough information to determine the matte. Initially, for each pixel p in Ω, F_p and B_p are approximated by corresponding the nearest foreground pixel in Ω_F and background pixel in Ω_B. Then, the constructed $(F-B)$ image is smoothed by a Gaussian filter to suppress significant changes due to noise and inaccurate estimation of F and B.

2. α **reconstruction** α is reconstructed by solving Poisson equations (6) using the current $(F-B)$ and ∇I.

3. F,B **refinement** Let $\Omega_F^+ = \{p \in \Omega | \alpha_p > 0.95, I_p \approx F_p\}$. The condition $\alpha_p > 0.95$ and $I_p \approx F_p$ guarantee that the pixels in Ω_F^+ are mostly foreground. Similarly, let $\Omega_B^+ = \{p \in \Omega | \alpha_p < 0.05, I_p \approx B_p\}$. Here, F_p, B_p and I_p represent the color vectors at pixel p. We update F_p and B_p according to the color of the nearest pixels in $\Omega_F \cup \Omega_F^+$ and in $\Omega_B \cup \Omega_B^+$, respectively. A Gaussian filter is also applied to smooth $(F-B)$.

We iterate the above steps 2 and 3 until change in the matting results is sufficiently small or both Ω_F^+ and Ω_B^+ are empty in step 3. Typically, only a few iterations are needed. In each iteration, the selection of Ω_F^+ and Ω_B^+ has little error, which guarantees that more accurate colors in these two regions are further propagated into less accurate neighboring pixels.

Global Poisson matting works well in scenes with a smooth foreground and background. However, for a complex image, Equation (3) may not be a good approximation of the matte gradient globally, where the background and foreground gradients cannot be ignored. In the next section, we bring the user into the matting loop to locally refine the global Poisson matting result.

4 Local Poisson matting

Equation (2) can be rewritten as:

$$\nabla\alpha = A(\nabla I - \mathbf{D}) \quad (7)$$

where $A = \frac{1}{F-B}$ and $\mathbf{D} = [\alpha\nabla F + (1-\alpha)\nabla B]$. A affects the matte gradient scale in that increasing A would sharpen boundaries. $\mathbf{D}$ is a gradient field caused by the background and foreground. Hence, we need to estimate A and $\mathbf{D}$ to approach the ground truth, A^* and $\mathbf{D}^*$. In *global Poisson matting*, A is automatically estimated from the image and $\mathbf{D}$ is assumed to be zero. When the background or foreground have strong gradients, global Poisson matting results in a poor quality mattes.

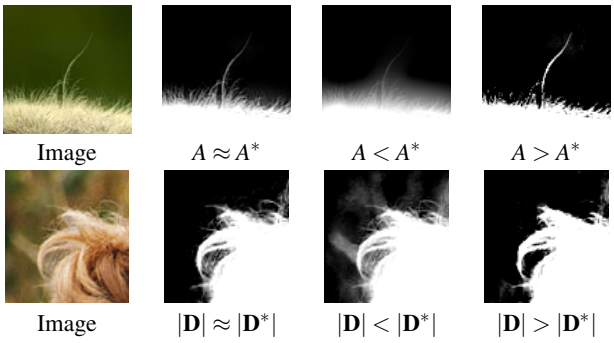


Figure 3: The mattes solved by Poisson matting when A and D approximate A^* and $\mathbf{D}^*$ differently. *Top*: $A \approx A^*$ results in a correct matte. $A < A^*$ results in a smooth matte and $A > A^*$ results in a sharp matte. *Bottom*: Similarly, $|\mathbf{D}| < |\mathbf{D}^*|$ or $|\mathbf{D}| > |\mathbf{D}^*|$ results in erroneous mattes.

In this section, we introduce *local Poisson matting* to allow users to locally manipulate the gradient field. Figure 3 shows the mattes generated with different A and $\mathbf{D}$. If A is smaller than A^*, the matte becomes smoother. Similarly, when $|\mathbf{D}|$ is different from $|\mathbf{D}^*|$, we obtain erroneous mattes. A key observation in local Poisson matting is that the user may be able to inspect the recovered matte and propose how to manipulate A and $\mathbf{D}$ to improve the matting result.

4.1 Poisson Matting in Local Region

To refine the result of *global Poisson Matting*, the user can specify a region Ω_L that s/he is not satisfied with, and apply local Poisson matting. As shown in Figure 2(b), the integral region becomes $\Omega_L \cap \Omega$ and the boundary of the new integral region becomes $\partial\Omega = \{p \in \overline{(\Omega_L \cap \Omega)} | N_p \cap (\Omega_L \cap \Omega) \neq \emptyset\}$. Figure 2(b) illustrates the user selection Ω_L and new boundary $\partial\Omega$. The variational problem to be minimized in local Poisson matting is given by:

$$\alpha^* = \arg\min_{\alpha} \int\int_{p \in \Omega_L \cap \Omega} ||\nabla\alpha_p - A_p(\nabla I_p - \mathbf{D}_p)||^2 dp. \quad (8)$$

with Dirichlet boundary condition $\alpha|_{\partial\Omega} = \widehat{\alpha}|_{\partial\Omega}$. The local Dirichlet boundary condition $\widehat{\alpha}|_{\partial\Omega}$ is defined as:

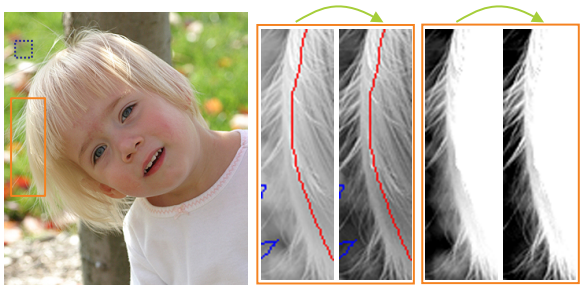
$$\widehat{\alpha}_p|_{\partial\Omega} = \begin{cases} 1 & p \in \Omega_F \\ 0 & p \in \Omega_B \\ \alpha_g & p \in \Omega \end{cases} \quad (9)$$

where α_g is the current matte value in an unknown region on the local boundary.

Usually, the local region size we use is small (fewer than 200×200 pixels), where a Poisson solver very quickly generates the matting result. Moreover, because of existing boundary conditions, the local operation is seamlessly propagated into the matte where no obvious boundary discontinuity can be seen.

4.2 Local Operations

Users can modify A and $\mathbf{D}$ in the selected region to produce a better approximation of $\nabla\alpha$. Two kinds of operations are provided: *channel selection* and *local filtering*. Channel selection reduces the error of $\mathbf{D}$ and local filtering manipulates A and $\mathbf{D}$ directly. By using these operations, users do not need to optimize matte in a pixel-wise manner, and the results after modification are quickly produced.



(a) (b) (c)

Figure 4: Channel selection. (a) The input image. The user selects samples from the blue region (dashed rectangle) in the background. (b) Images in grayscale and optimized channels, where $\partial\Omega$ is shown as a red/blue line. (c) Alpha mattes computed from grayscale and the optimized channels. A clearer matte is generated from the optimized channel.

4.2.1 Channel Selection

For color images, Equation (3) can be measured in different channels, e.g. any single R/G/B or grayscale channel. In global Poisson matting, we have used the grayscale channel, $\nabla\alpha = A_g(\nabla I_g - \mathbf{D}_g)$. The smooth background or foreground assumption makes $|D_g|$ small such that $A_g\nabla I_g$ is a good approximation of $\nabla\alpha$. Similarly, in a local region, we try to construct a new channel $\gamma = aR + bG + cB$ with a smooth background or foreground in which $|\mathbf{D}_\gamma|$ is smaller than $|\mathbf{D}_g|$. Thus, we minimize the variance of the foreground or background colors, and the new channel γ is constructed as follows:

1. Users select either background or foreground color samples $\{(R_i\ G_i\ B_i)\}_{i=1}^N$ in the image using a simple brush. In the accompanying video, we show an example where background samples are used for channel selection.

2. Compute the weights $(a\ b\ c)$ to minimize the sample variances $\sum_i^N(\gamma_i - \overline{\gamma})^2$ in the new channel. This is a linearly constrained quadratic optimization problem:

$$\min_{a,b,c}\sum_i[(a\ b\ c)\cdot(R_i\ G_i\ B_i)^T - (a\ b\ c)\cdot(\overline{R}\ \overline{G}\ \overline{B})^T]^2 \ \text{ s.t. } \ a+b+c=1 \tag{10}$$

where $(\overline{R}\ \overline{G}\ \overline{B})$ is the mean color value of the samples. The weights $(a\ b\ c)$ are obtained by solving an augmented linear system [P. Gill and Wright 1981]. Figure 4 shows the matte improvement in the new optimized channel as compared to the gray channel. The error of $\mathbf{D}$ is reduced and the hair shape is better recovered.

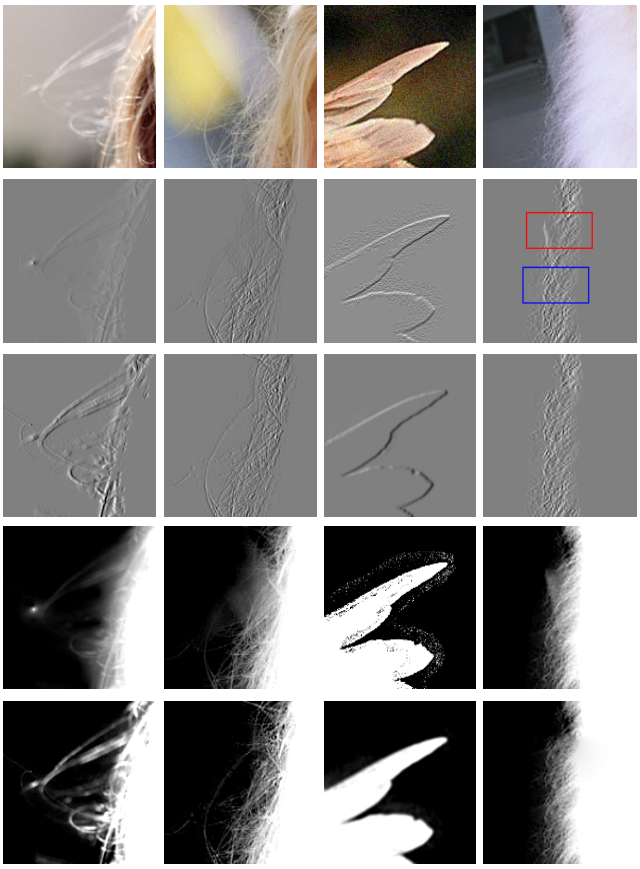
4.2.2 Local Filtering

To directly operate on A and $\mathbf{D}$, we also provide several local filters for users to manipulate the matte gradient field.

Boosting brush. When the matting result is smoother or sharper than what users expect, a *boosting brush* can be used to increase or decrease A directly. The boosting brush has a local Gaussian shape for each pixel p in the brush area. A_p is modified to A'_p using the boosting brush

$$A'_p = [1 + \lambda\exp(-\frac{||p-p_0||^2}{2\sigma^2})]\cdot A_p \tag{11}$$

where p_0 is the coordinate of the brush center, σ and λ are user defined parameters to control the size and strength of the boosting effect. Hence, the user can boost the whole or partial region by using brushes of various sizes. If $\lambda > 0$, this filter will increase A around the brush center, and vice versa. Figure 5 (a) shows that a smooth alpha matte is sharpened after applying the boosting brush.



(a) Boosting (b) Highpass (c) Diffusion (d) Clone

Figure 5: Local filtering. The top row shows input images. The second and third rows are approximated matte gradients (x-direction only) before and after applying local filters. The bottom two rows are corresponding alpha mattes computed from local Poisson matting. (a) Boosting brush produces sharper mattes. (b) Highpass filter recovers structures of matte. (c) Diffusion filter removes noise. (d) Clone brush copies the matte gradient from the blue region and pastes it into the red region.

Highpass filtering. The channel selection operation generates a smooth background or foreground, leading to low frequency background or foreground gradients. Therefore, D can be estimated using the low-frequency part of the image gradient:

$$\mathbf{D} = K * \nabla I$$

where $K = N(p; p_0, \sigma^2)$ is a Gaussian filter centered at pixel p_0 and $*$ is the convolution operator. From Equation (7), $\nabla\alpha = A(\nabla I - K * \nabla I)$, where $(\nabla I - K * \nabla I)$ corresponds to a highpass filter. Figure 5(b) demonstrates that a clear alpha matte structure is recovered after applying highpass filtering.

Diffusion filtering. On the boundary of a solid object where alpha matte changes quickly, $A\nabla I$ is already a good approximation [Mitsunaga et al. 1995]. However, the image gradient ∇I is sensitive to noise and blocking effects in JPEG images. We adopt anisotropic diffusion [Perona and Malik. 1990] to diffuse the image. It is an edge-preserving blurring process to remove small scale noise. Then the image gradient ∇I is re-computed from the diffused image. Figure 5(c) shows that the noise is suppressed in the diffused image.

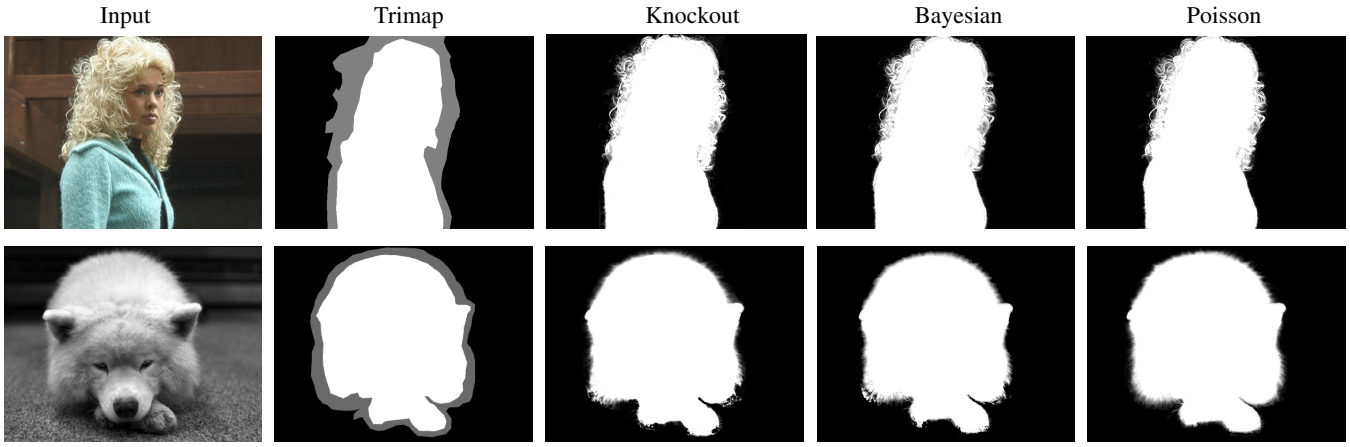


Figure 6: Global Poisson matting results comparison. Two result mattes are shown. For the image in the upper row ©[Chuang et al. 2002], comparable results are generated. In the lower row ©Philip Greenspun, our result has less visibel artifacts.

Clone brush. In some difficult situations, the clone brush can be used to directly copy the matte gradient $A(\nabla I - \mathbf{D})$ from a user selected source to a target region. In Figure 5(d), the matte gradients in the blue region are selected for pasting into the red region. Note that the cloned matte gradient produces a "convincing" matte in the situation that the foreground and background are almost indistinguishable.

Another two brushes are also helpful. One is *Erase brush* that removes unwanted alpha matte directly. Another is *Inverse brush* that inverses the incorrect sign of matte gradient. We demonstrate the filtering effects in the matte gradient and the corresponding alpha matte in Figure 5, which serves different purposes in local processing.

4.2.3 Refinement Process

With the above local operations, we allow users to refine the matte in a selected region. Based on the global Poisson matting result, local Poisson matting proceeds as follows:

1. Apply channel selection to reduce the errors in $\mathbf{D}$. For solid object boundaries, apply the diffusion filter to remove possible noise.
2. Apply the highpass filtering to obtain an approximation of $\mathbf{D}$.
3. Apply the boosting brush to manipulate A.
4. Possibly apply the clone brush if gradients are indistinguishable.

The erase brush and inverse brush can be optionally applied in any step. At each step, local Poisson matting can generate results very quickly. The user can observe the matting result and select any remaining unsatisfactory regions for further refinement.

5 Results and Applications

We apply Poisson matting to many complex images to demonstrate the effectiveness of our approach, and extend it to two applications.

We compare the constructed mattes from our global Poisson matting in Figure 6 with those from Corel Knockout®and our implementation of Bayesian matting, given the same trimap. In the first row, we use the image that appeared in [Chuang et al. 2002], where complex hair opacity exists in simple foreground and background colors. Our method generates comparable results. In the second row, our method reduces abrupt matte gradient changes because of the local continuity property of the solution of Poisson equations. As a result, a better matte with little visual artifact is constructed. On the 640×480-pixel image in Figure 6, our method took 1.3 and 1.1 seconds to converge for the "girl" and "Samoyed dog" image respectively using the trimap shown; Knockout took 0.9 and 0.8 seconds, and our implementation of the Bayesian method took 28.8 and 22.7 seconds.

Figure 9 shows the results of natural image matting using local Poisson matting. Most images are very complex not only in foreground and background colors, but also in their gradient fields. Therefore, Bayesian matting which relies on sampling color pixels does not work well. Our local Poisson matting helps users to conveniently improve matting results. According to our experience, the average time to process a 600×400-pixel image in Figure 9 is less than 10 minutes. For the "dog" example in the third row, fewer than 7 local refinements are sufficient. For the images in the second and fourth rows, we perform no more than 20 local refinements of various region sizes.

By closely inspecting the input images and matting results, we have found that Bayesian matting fails in places where foreground and background colors are similar, or their colors have large change so that correct samples are buried in the local region. To see the difference clearly, we zoom in these problematic local regions in the fourth column of Figure 9.

Multi-background. Our method can be applied to matting with multiple backgrounds as shown in Figure 7. The triangulation solution [Smith and Blinn 1996] has similar settings; however, it requires known backgrounds. Let $\{I_t\}_{t=1}^T$ be images of the same foreground F with multiple different backgrounds $\{B_t\}_{t=1}^T$. Without any information about backgrounds, we calculate the mean image

$$\overline{I} = \frac{1}{T}\sum_{t}^{T}(\alpha F + (1-\alpha)B_t) = \alpha F + (1-\alpha)\overline{B} \qquad (12)$$

where $\overline{B} = \frac{1}{T}\sum_t^T B_t$ is the mean background. Therefore, $\overline{B}$ is usually a smoothed image for all backgrounds in $\{B_t\}_{t=1}^T$. Poisson matting works better in the mean image than in any individual image in $\{I_t\}_{t=1}^T$. An example is shown in Figure 7 where global Poisson matting is used, followed by local Poisson matting for a region at the lower right part of the image.

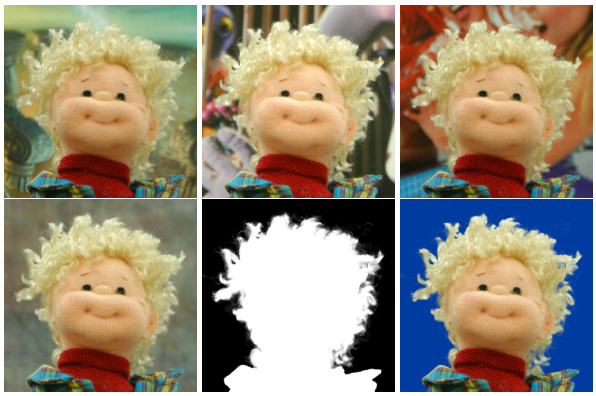


Figure 7: Multi-Background. Top row: three input images with different complex backgrounds. Bottom row: the mean of all eight input images, computed alpha matte and composite image using Poisson matting.



Figure 8: De-fogging. The de-fogged image is obtained using the boosting brush with several local operations.

De-fogging. A simple fog model introduced in [Narasimhan and Nayar 2003] can also be represented as a matting equation:

$$I = I_c \cdot e^{-\beta d} + Fog \cdot (1 - e^{-\beta d}) \tag{13}$$

where Fog is the color of fog, I_c is the clear image without fog, β is the scattering coefficient of the atmosphere and d is the depth value. [Narasimhan and Nayar 2003] recovered the clear image I_c by manually assigning depth, scattering coefficient and fog color information in a global manner. We provide a de-fogging method that can locally assign the information to produce good results. After taking the partial derivatives of (13), we obtain $\nabla I_c = \nabla I \cdot e^{\beta d}$ with respect to assigned depth d. We use the boosting brush proposed in Section 4.2.2 to locally modify $e^{\beta d}$ in the selected region to fine tune the de-fogged image gradient. By solving Poisson equations, any modification to the image can be automatically and seamlessly propagated to generate the de-fogged image without any visible boundaries. An example is shown in Figure 8.

6 Conclusions

In this paper, we have presented a new digital matting method, Poisson matting. By solving Poisson equations, Poisson matting reconstructs a faithful matte from its approximated gradient field estimated from an input image semi-automatically. Given a few hints using local operations, Poisson matting is capable of producing impressive results for many complex images problematic to previous natural image matting methods.

The matting problem is under-constrained and intrinsically difficult. While we believe Poisson matting has made some important advancements in this problem, some limitations remain to be addressed. First, when the foreground and background colors are very similar, the matting equation becomes ill-conditioned. In this case, the underlying structure of the matte can not be easily distinguished from noise, background or foreground. The second difficulty arises when the matte gradient estimated in global Poisson matting largely biases the true values, so that small regions need to be processed for local refinements in local Poisson matting, which increases user interaction. Last, when the matte gradients are highly interweaved with the gradients of the foreground and background within a very small region (e.g., two human figures with long hair that sightly overlap in the image). Effective user interaction is an issue in this difficult situation.

In the future, we are interested in combining Poisson matting with the Bayesian method and extending our work to video matting.

Acknowledgements. We would like to thank the anonymous reviewers for their constructive critiques. Many thanks to Stephen Lin for his help to improve the manuscript. Chi-Keung Tang's research is supported in part by the Research Grant Council of Hong Kong Special Administration Region, China: HKUST6171/03E and AOE/E-01/99.

References

BERMAN, A., VLAHOS, P., AND DADOURIAN, A. 2000. Comprehensive method for removing from an image the background surrounding a selected object. *U.S. Patent 6,134,345*.

CHUANG, Y.-Y., CURLESS, B., SALESIN, D. H., AND SZELISKI, R. 2001. A bayesian approach to digital matting. In *Proceedings of CVPR 2001, Vol. II, 264-271*.

CHUANG, Y.-Y., AGARWALA, A., CURLESS, B., SALESIN, D. H., AND SZELISKI, R. 2002. Video matting of complex scenes. In *Proceedings of ACM SIGGRAPH 2002*, 243–248.

ELDER, J. H., AND GOLDBERG, R. M. 2001. Image editing in the contour domain. *IEEE Trans. Pattern Anal. Machine Intell. 23(3): 291-296*.

FATTAL, R., LISCHINSKI, D., AND WERMAN, M. 2002. Gradient domain high dynamic range compression. In *Proceedings of ACM SIGGRAPH 2002*, 249–256.

FINLAYSON, G. D., HORDLEY, S. D., AND DREW, M. S. 2002. Removing shadows from images. In *Proceedings of ECCV 2002, Vol. IV, 823-836*.

HILLMAN, P., HANNAH, J., AND RENSHAW, D. 2001. Alpha channel estimation in high resolution images and image sequences. In *Proceedings of CVPR 2001, Vol. I, 1063-1068.*

MITSUNAGA, T., YOKOYAMA, T., AND TOTSUKA, T. 1995. Autokey: Human assisted key extraction. In *Proceedings of ACM SIGGRAPH 1995*, 265–272.

NARASIMHAN, S., AND NAYAR, S. 2003. Interactive deweathering of an image using physical models. *IEEE Workshop on Color and Photometric Methods in Computer Vision 2003*.

P. GILL, W. M., AND WRIGHT, M. 1981. Practical optimization. *Academic Press, Boston, MA, USA*.

PÉREZ, P., GANGNET, M., AND BLAKE, A. 2003. Poisson image editing. In *Proceedings of ACM SIGGRAPH 2003*, 313–318.

PERONA, P., AND MALIK., J. 1990. Scale space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Machine Intell. 12(7):629-639*.

QIAN, R. J., AND SEZAN, M. I. 1999. Video background replacement without a blue screen. In *Proceedings of ICIP 1999*, 143–146.

RUZON, M. A., AND TOMASI, C. 2000. Alpha estimation in natural images. In *Proceedings of CVPR 2000*, 18–25.

SMITH, A. R., AND BLINN, J. F. 1996. Blue screen matting. In *Proceedings of ACM SIGGRAPH 1996*, 259–268.



Figure 9: Local Poisson matting results. The first column shows input images. The middle three columns compare the mattes generated by Bayesian matting and our method. In the last column, we composite our extracted foreground on other background images.