Actively Illuminated Objects using Graph-Cuts

Asla Sá
Visgraf/IMPA
Estrada Dona Castorina 110
Rio de Janeiro, Brazil
asla@visgraf.impa.br

Marcelo Bernardes Vieira
DCC - ICE - UFJF
Cidade Universitária
Juiz de Fora, MG, Brazil
marcelo.bernardes@ufjf.edu.br

Anselmo Antunes Montenegro
IME UERJ
R. São Francisco Xavier 524
Rio de Janeiro, Brazil
anselmo@ime.uerj.br

Paulo Cezar Carvalho
Visgraf/IMPA
pcezar@visgraf.impa.br

Luiz Velho
Visgraf/IMPA
lvelho@visgraf.impa.br

Abstract

This paper addresses the problem of foreground extraction using active illumination and graph-cut optimization. Our approach starts by detecting image regions that are likely to belong to foreground objects. These regions are constituted by pixels where the difference in luminance for two differently illuminated images is large. The foreground objects are segmented by graph-cut optimization using those regions as a seed and using a energy function based on probability distributions derived from both input images and their difference. Several light sources and different illumination schemes can be used to mark the foreground. Our method has only two scalar parameters which can be set once for a wide variety of scenes.

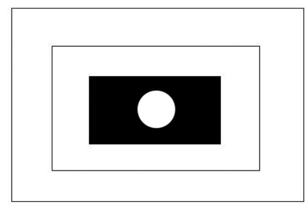

Figure 1. Ambiguity: the decision on whether the central circle is background or foreground is a matter of interpretation.

1 Introduction

Image segmentation is an important problem in Computer Vision. A special case of the general segmentation problem is the foreground/background image segmentation, in which a binary classification is applied to an image that has a perceptual background/foreground separation.

Given a single image, some additional knowledge has to be given to the solver as an initial clue of what is background or foreground. Figure 1 shows the intrinsic ambiguity in determining the foreground for a general image. Other examples of ambiguity are scenes with more then two distinct planes of information, where the classification of each plane as background or foreground is a matter of interpretation.

In practice, the information is usually disambiguated either by user interaction or by the acquisition of additional information about the scene such as previous calibration of the background, defocusing, analysis of movement and other techniques most of them benefiting from analyzing a collection of images.

In this work we propose the use of active lighting to insert a priori clues to solve the segmentation problem [22]. In order to do that, we introduce a light source that illuminates the foreground objects to be segmented more strongly than the background. The information derived from this difference in illumination replaces the indication of object and background pixels by the user. These initial clues are then used as seeds for a graph-cut minimization algorithm to obtain a high quality segmentation.

The paper is organized as follows: previous work is presented in the next section; in section 3 the segmentation problem is addressed and our active approach to the problem is presented; in section 4 the objective function is derived; in section 5 we present the segmentation results while conclusions and future work are discussed in the final section.

2 Related work

2.1 Active lighting

is a familiar tool to the 3D photography community. The most commonly adopted setup uses a projector as the light source that communicates with the camera, as in [22].

In recent works, camera flash was explored to enhance image quality. [6], [17] have proposed the use of bilateral filter to decompose a flash/non-flash pair of images and then recombine them appropriately; both authors have to deal with shadows produced by the flash. Multiple images with flash positioned in different locations among them is used to extract object borders in [18] and a non-photo realistic rendering is then applied to the images. Our work differs from these previous works since it explores a light source, especially positioned to illuminate only the foreground objects to be segmented, that stays in a fixed position between shots; while its intensity is modulated, but not turned off, what do not produces shadows inconsistencies.

2.2 Foreground segmentation

There are two main branches of research in this area: one branch assumes that image acquisition can be controlled to produce images that are automatically segmented, while the other analyzes a given image without any assumption about image formation. In the first approach the clues are decided *a priori* and the segmentation is automatic whereas in the latter the initial clues for solving the problem are inserted by the user *a posteriori*.

An example of a priori information widely adopted is chroma key [24]. Recently much work has been done in proposing segmentation methods that defines clues a posteriori intending to minimize user intervention. In most cases the user has to indicate coarsely the foreground and the background pixels [19],[25] as initial restrictions for a minimization process.

2.3 Image segmentation by graph cuts

The problem of image segmentation is a special case of a pixel labeling problem. It can be modeled as an optimization problem, which consists of computing the best image segmentation among all possible image segmentations satisfying a set of predefined restrictions.

Many approaches found in the literature model the problem of image segmentation as a problem that requires the minimization of certain cost functions. Some of these methods, for example, those based on *shortest paths* [15] and *snakes* [11], [5] are limited to 2D applications, where the segmentation boundary is determined by a 1D curve. Other approaches, as those based on *level set methods* [16] and *normalized cuts* [23], compute only approximated solutions, which can be arbitrarily far from the global optimum.

Recently, Graph-Cut Minimization became widely used in image segmentation [1], [19], [4]. In the next sections, we briefly review the basis for these methods.

2.4 Energy minimization in vision

One of the most usual applications of energy minimization in Computer Vision is to solve *pixel labeling problems*, which generalize problems as stereo and motion, image restoration and segmentation.

In pixel labeling problems the goal is to find a labeling $f : \mathcal{P} \mapsto \mathcal{L}$, that maps a set of pixels $\mathcal{P}$ to a set of labels $\mathcal{L}$, which minimizes some energy function. This energy function typically has the form

$$E(f) = \sum_{p \in \mathcal{P}} D_p(f_p) + \sum_{p,q \in \mathcal{N}} V_{p,q}(f_p, f_q), \tag{1}$$

where $\mathcal{N} \in \mathcal{P} \times \mathcal{P}$ is a neighborhood system on $\mathcal{P}$. $D_p(f_p)$ is a function based on the observed data that measures the cost of assigning label f_p to p and $V_{p,q}(f_p, f_q)$ is an spatial smoothness term that measures the cost of assigning labels f_p and f_q to adjacent pixels p and q.

Energy functions like E are, in general, very difficult to minimize, as they are non convex functions in large dimension spaces. When these energy functions have special characteristics, it is possible to find their exact minimum using dynamic programming. Nevertheless, in the general case, it is usually necessary to rely on general minimization techniques, like Simulated Annealing [7], which can be very slow in practice.

An interesting property of a graph cut C is that it can be related to a labeling f, mapping the set of vertices $\mathcal{V} - \{s, t\}$ of a graph $\mathcal{G}$ to the set $\{0, 1\}$, where $f(v) = 0$, if $v \in S$, and $f(v) = 1$, if $v \in T$. This labeling defines a *binary partitioning* of the vertices of the graph.

2.5 Energy minimization by graph cuts

Graph cut optimization was proposed as an efficient way to minimize a larger class of energy functions than those that can be minimized by dynamic programming. It has become very popular in the area of early vision, and has been successfully used to solve many different problems as image restoration [4], [3],[8], stereo and motion [4], [3], [9], [20], [21], image synthesis [14], multi camera scene reconstruction [12] and image segmentation [10], [2].

The optimality of graph cut minimization methods, considering labeling problems, depends on the number of labels and the exact form of the smoothness term V. In [8] is proved that the method yields global minimum solutions when the problem is a binary labeling problem, while in

[9] proved that, for labeling problems with arbitrary number of labels, it is possible to compute global minimal, if the smoothness term is restricted to a convex function.

In most cases, it is necessary to preserve boundary discontinuities, which is not possible by using a convex function as the smoothness penalty term in the energy function. In general, the minimization for energy functions that preserve discontinuities by graph cut minimization can only produce approximate solutions. In [4] Boykov et al proposed a graph cut based algorithm, named Expansion Move Algorithm, that is able to compute a local minimum for discontinuity preserving energy functions. They proved that this local minimum lies within a small multiplicative factor (equal to 2) of the global minimum.

The early proposals that used graph cut optimization as a method for energy minimization required the construction of a specific graph for each particular problem. In [13] is introduced a general scheme for graph cut minimization of energy functions that belong to the class of regular functions. In our work, the graph cut optimization algorithm used is based on such formulation.

3 Active illumination

Active illumination can be combined with graph-cut optimization to perform the segmentation of foreground regions. We call this *active segmentation*, meaning the use of a light sources to illuminate only the objects to be segmented leaving the background essentially unchanged. Thus, the light source works as a substitute to the user, acting on the scene to indicate object and background seed elements automatically. This pre-segmentation provides the color distribution of each region, that can be used in a graph-cut optimization step to obtain the final segmentation.

In Fig. 2 the active approach is illustrated. Fig. 3 illustrates another example where the character was actively illuminated. In this case the background is composed of some specular elements (the leaves) creating a difficult scenario to the non-optimized method (as can be observed in next section).

Figure 2. (upper-left) and (upper-right) are the input images differently illuminated by varying the camera flash intensity between shots, (lower) is the difference thresholded image.

Figure 3. Input images differently illuminated.

3.1 Graph-cut optimization

The initial seed obtained by the active approach described above gives us important clues about the regions that are likely to belong to the background and foreground of the scene. Based on these clues, it is possible to compute the desired segmentation by minimizing an energy function. If this energy function is chosen in such a way that some regularity properties hold, then it is possible to minimize it efficiently by graph cut optimization methods.

As in [3] and [2], the energy function that we use is a discontinuity preserving energy function, obtained in the context of a Maximum a Priori Markov Random Field estimation. It is defined in terms of a set of pixels $\mathcal{P}$, a set of pairs of neighboring pixels in a neighborhood system $\mathcal{N}$ and a binary vector $A = (A_1, A_2, .., A_{|p|})$, where A_p is the assignment of pixel p either to 0 (background) or 1 (foreground).

The energy function has the form of the following cost function:

$$E(A) = \lambda \sum_{p \in \mathcal{P}} R_p(A_p) + \sum_{\{p,q\} \in \mathcal{N}} B_{\{p,q\}} \cdot \delta(A_p, A_q) \quad (2)$$

This energy function is defined in terms of a *regional term*, that measures the fitness of a region to the background or foreground of the scene, and a *boundary term*, which penalizes discontinuities in the label assignment while preserving those that are associated to features of the image. The first term is a *regional term* that measures how the intensities of the pixels of the image fit into intensity models (for example, obtained by a histogram) of the background and foreground. The second term is a *boundary term*, which penalizes discontinuities in the label assignment while preserving those that are associated to features of the image. Coefficient $B_{\{p,q\}} > 0$ can be interpreted as a penalty for spatial discontinuity of the labels assigned to neighboring pixels p and q. $B_{\{p,q\}}$ should be large when pixels p and q are similar, and close to zero when p and q are very different, so that feature discontinuities are preserved. The $\lambda \geq 0$ constant is used to specify the relative importance of the regional term versus the boundary properties term.

The energy function proposed in our model is minimized by a graph cut optimization algorithm that follows the scheme proposed in [13]. Differently from [2], which is mainly interested in interactive segmentation, our approach does not rely on hard constraints. The regions determined by the active illumination thresholding are used as seeds to the graph cut optimization. However, their labels can be modified as the process is executed. Another important difference is that we work in the *Lab* color space.

In the next section we describe in detail the energy function that models the binary segmentation problem of energy minimization by graph cuts.

4 The objective function

The objective function is based on probability distributions of color values in three regions: background, foreground object and boundary. They are defined assuming the following:

- Most actively illuminated pixels belongs to the foreground objects. Note that the influence of active illumination on the background can lead to wrong overall segmentation;
- The actively illuminated regions capture the object features, that is, they contain all color information necessary to distinguish foreground objects from the background;
- Regions corresponding to moving objects in the scene represent a small fraction of the scene;
- Color differences in Lab space are sufficient to define relevant object/background boundaries.

The challenge is to define probability distributions that approximate the real distribution of the expected segmentation regions. For the background region, we employ *a priori* distributions of the luminance difference $L_{I_2} - L_{I_1}$. Color histograms from the seed regions are used to build a color distribution function for the foreground region.

These distributions, together with a boundary likelihood function based on distances in Lab space, are the basis of the cost function to be proposed.

4.1 Composing the cost functions

The goal is to find the labels $\mathbf{X} = \{x_p, \;\; p \in I_1\}$, where x_p is 0 if p is background or 1 if p is foreground, that minimize an objective function $E(\mathbf{X})$. The regional term of the energy function is defined as:

$$C(x_p) = \begin{cases} -\log(\mathrm{p}_O(p)), & \text{if } x_p \text{ is 1 (foreground)} \\ -\log(\mathrm{p}_B(p)), & \text{if } x_p \text{ is 0 (background)} \end{cases} \tag{3}$$

The boundary term, for neighboring pixels p, q is $-|x_p - x_q| \log \mathrm{p}_R(p,q)$. Thus the final objective function is

$$E(\mathbf{X}, \sigma_L, \sigma_C) = \sum_{p \,\in\, I_1} C(x_p) - \sum_{p,q \,\in\, I_1} |x_p - x_q| \cdot \log \mathrm{p}_R(p,q), \tag{4}$$

where points q are those in the 8-connected neighborhood of p.

We turn now to the definition of $\mathrm{p}_B(p)$, $\mathrm{p}_O(p)$ and $\mathrm{p}_R(p,q)$. We start by discussing how to infer foreground sites from the input data. With the above assumptions, high values of the luminance difference $|L_{I_2}(p) - L_{I_1}(p)|$ indicate foreground pixels p, where $L_{I_1}(p)$ and $L_{I_2}(p)$ are the luminance channels of the transformed images I_1 and I_2. However, it cannot be stated that low values of that difference indicate background pixels since there may be parts of foreground objects that are not actively illuminated (like shadow areas). Thus, the luminance difference does not characterize completely foreground and background elements.

Luminance difference for background pixels can be modeled by a gaussian distribution, with density

$$\mathrm{p}_B(p) = \frac{1}{\sqrt{2\pi}\sigma_L} \exp\left(\frac{-|L_{I_2}(p) - L_{I_1}(p)|^2}{2\sigma_L^2}\right), \tag{5}$$

where σ_L is the standard deviation of the luminance differences, illustrated in Fig. 4.

High $\mathrm{p}_B(p)$ values do not necessarily indicate that p is background but pixels with small $\mathrm{p}_B(p)$ values are likely to belong to the foreground. The set of foreground pixels are then defined as $O = \{p \;\mid\; \mathrm{p}_B(p) \;<\; t\}$, where t is a small threshold. We fix $t = 0.05$ since the parameter σ_L can be adjusted.

The color histogram of the foreground pixels determine the object probability function. For simplicity, we use a 3D histogram for the Lab components with uniform partition.

Figure 4. Images of background probabilities. Darker pixels have smaller probabilities (left) $\sigma_L = 15$ **(center)** $\sigma_L = 25$ **(right)** $\sigma_L = 35$**.**

Let nb_L, nb_a and nb_b be the number of predefined bins for each lab component. All points $p \in O$, with normalized color components $L_1(p)$, $a_1(p)$ and $b_1(p)$, are assigned to a bin k with coordinates

$$(\lfloor L_1(p) * nb_L) \rfloor, \lfloor a_1(p) * nb_a \rfloor, \lfloor b_1(p) * nb_b \rfloor)$$

The object distribution function is then defined as

$$\mathrm{p}_O(p) = \frac{n_k}{n_O} \qquad (6)$$

where n_k is the number of pixels assigned to the bin k and n_O is the number of pixels in the object region O, illustrated in Fig. 5.

To construct the histogram information only one image is considered and it will depend on each situation. In our experiments $L_1(p)$, $a_1(p)$ and $b_1(p)$ are the color components from the image correspondent to the lowest projected ρ value. Another consideration is that it may be difficult to determine the number of bins for each component. The bins that distinguish relevant color groups when the partition is uniform. If the number of bins is too small, wide ranges are mapped in few bins. If there are too many bins, frequencies tend to be small everywhere. In our experiments, the object pixels are sufficient to populate a histogram with n_L=32, n_a=64 and n_b=64 bins.

Figure 5. Object probabilities: darker pixels represents smaller probabilities. (left)$\sigma_L = 15$**, (center)**$\sigma_L = 25$ **and (right)**$\sigma_L = 35$**.**

Finally, the likelihood function for neighboring boundary pixels is given by

$$\mathrm{p}_R(p, q) = 1 \; - \; \exp(\frac{-(||Lab(p) - Lab(q)||)^2}{2\sigma_C^2}), \qquad (7)$$

where $Lab(p)$ denotes the color at point p and σ_C is the standard deviation of the L^2-norm of the color difference, illustrated in Fig. 6. The effect of this term is that, if the colors of pixels p and q are close in the Lab space, their connection is unlikely to cross a foreground-background border.

Figure 6. Image of boundary probabilities taking the maximum value of a 8-connected neighborhood. Darker pixels have smaller probabilities (left) $\sigma_C = 5$ **(center)** $\sigma_C = 15$ **(right)** $\sigma_C = 25$**.**

According to [13], an energy function ot the form

$$E(x_1, x_2, \cdots, x_n) = \sum_p E^p(x_p) + \sum_{p \neq q} E^{p,q}(x_p, x_q)$$

, where each x_p is a 0-1 variable, can be minimized by means of a minimum graph-cut when it is regular, that is, satisfies the inequality

$$E^{p,q}(0,0) + E^{p,q}(1,1) \leq E^{p,q}(0,1) + E^{p,q}(1,0).$$

In our case, when $x_p = x_q$, we have $E^{p,q}(0,0) = E^{p,q}(1,1) = E^{p,q}(0,1) = E^{p,q}(1,0) = 0$. On the other hand, if $x_p \neq x_q$, then $E^{p,q}(0,1) = E^{p,q}(1,0) = 0$ and $E^{p,q}(0,1) = E^{p,q}(1,0) = -\log(p_R(p,q) \geq 0$, since $0 \leq p_R(p,q) \leq 1$. Hence, the proposed energy function is regular.

5 Method and results

The main steps of our active segmentation method are illustrated in Fig. 7. Two input images are acquired as described in the following section. We apply a low-pass filter to the input images in order to reduce noise. Next, the input colors are transformed into the Lab color system. This perception-based color space is desirable for two reasons: we need to cluster regions with small perceived color variations and we want to explore the orthogonality between luminance and chromaticity information in Lab space. Our goal is to have perceptually homogeneous regions, with the segmentation boundaries preferably located where high Lab color differences occur.

Active illumination is explored to attribute weights to the pixels that are used in the energy minimization by graph cuts. For the optimization step, a graph where the nodes are

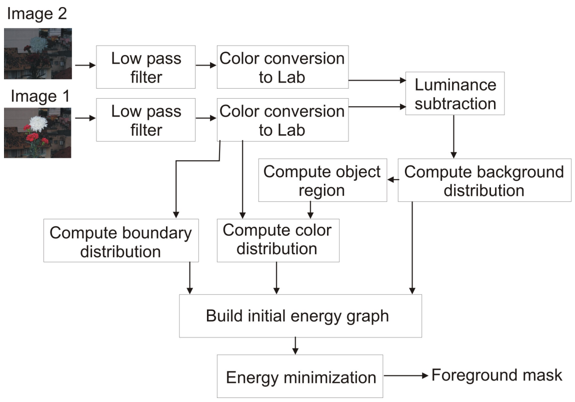

Figure 7. The proposed active foreground extraction method.

the pixels and the edges form a 8-connected neighborhood is created. The object and background color distributions are used to compute the cost of assigning a object or background label to each node. The boundary distribution is used to compute the cost of having an object/background transition for each edge.

Initially, all points belonging to the estimated object region are labeled as foreground. All other points are labeled as background. A min-cut/max-flow algorithm is used to find the global minimum

$$\widehat{\mathbf{X}} = \text{arg min}_{\mathbf{X}} \; E(\mathbf{X}, \sigma_L, \sigma_C) \tag{8}$$

Only some object regions can be determined before the optimization step. Usually, seeds for both the foreground and the background are used, which implies that histograms for both classes are available. The seed pixels are not allowed to have their label changed. In our case, there is no guarantee that the estimated object region is correct. Furthermore, the *a priori* background distribution (Eq. 5) is not fully precise. Equation 8 is then defined in such a way that the original labels may be changed during optimization.

A modified version of the energy minimization software was used in our implementation. Basically, the constraint that the original seeds must be kept was removed. The result is an image mask for foreground pixels.

The parameter σ_L determines the number of pixels in the initial object region. Lower values result in more pixels as shown in Fig. 4. Depending on the object material, it is remarkable that even small variations of the illumination can be detected by luminance differences. On the other hand, objects with highly specular, transparent or with complex structure materials are hard to be detected. This is the case of the hair in the images of Fig. 3. The object membership probabilities in the hair region (Fig. 5) show how the number of initial pixels affects the color distribution.

Figure 8. Segmentation results of the images in Fig. 3 for several values of σ_L with $\sigma_C = 10$. Compare with Fig. 4. (left) $\sigma_L = 15$ (center) $\sigma_L = 25$ (right) $\sigma_L = 35$.

Fig. 8 shows the segmentation results for several values of $\sigma_L \in [15, 20]$. It is harder to segment the hair as σ_L assumes higher values. This happens when the pixels with high probability of this region are not enough to populate the histogram.

Figure 9. Segmentation results for the images in Fig. 3 for several values of σ_C with $\sigma_L = 20$. Compare with Fig. 6. (left) $\sigma_C = 5$ (center) $\sigma_C = 15$ (right) $\sigma_C = 25$.

Parameter σ_C controls how the image borders constrain the expansion or contraction of the object clusters during optimization. If its value is low, the difference of probability between the highest and lowest gradient values is high. As a result, the segmentation tends to be more fragmented and well aligned with high color variation areas. If the value of σ_C is high, the probabilities tend to vary more slowly, resulting in a smoother segmentation.

Segmentation results varying σ_C are shown in Fig. 9. Note that small background regions appear like holes inside the hair when $\sigma_C = 5$. Higher values of σ_C tend to classify these holes as foreground. As explained above, high values of σ_C smooth the segmented clusters.

In Figure 10 the segmentation result related to the input images shown in Figure 2 are presented. Observe the difference in the segmentation continuity when we vary the system parameters. In Figure 11 another example of input images and its segmentation is illustrated.

6 Conclusions

In this paper we introduce a new method for the segmentation of a scene into foreground and background regions.

Figure 10. (upper)Boundary probabilities with $\sigma_C = 5$. The resultant segmentation using different parameters: (lower-left) $\sigma_L = 25$, $\sigma_C = 10$, (lower-right) $\sigma_L = 15$, $\sigma_C = 5$.

Figure 11. Input images and the final segmentation mask.

The method is based on active illumination and employs graph cut optimization to segment the image.

The key idea exploited in our method is that light variations can be designed to affect objects that are closer to the camera. In this way, a scene is lit with two different intensities of an additional light source that we call *segmentation light source*. By capturing a pair of images with such illumination, we are able to distinguish between objects in the foreground and the scene background.

The main technical contributions of this work are the concept of foreground / background segmentation by active lighting and the design of a suitable energy function to be used in the graph cut minimization.

The method is fully automatic and does not require user intervention to label the image. Moreover, the energy function has only two parameters that must be specified: the standard-deviations of the normal background and boundary distributions. These parameters can be tuned only once for a wide variety of images with similar light setup.

Concerning the operating range of the method, we point out that the active light source positioning is crucial to the method. It can be positioned in order to illuminate only the object to be segmented without affect the background. The discussion of how far the background should be for the method to work is relevant only when the light source is affecting the background. In this case, background distance is dependent of sensors sensitivity, usually 3 stops of light intensity decay is enough for our method to segment correctly.

The quality of the masks produced by our method is, in general, quite good. Some difficult cases may arise when the objects are highly specular, translucent or have very low reflectance. Because of its characteristics, this method is suited for applications in which the user can control the scene illumination, for example in studio situations and/or using a flash/no-flash setup.

We remark that the additional information resulting from capturing two images of the same scene can be used also for other purposes, such as extending the dynamic range of the image and color correction.

This method can be naturally extended to active segmentation of video sequences. All that it is required for this purpose is a synchronized camera/projector system. We have already implemented the active segmentation for video [22] without performing the optimization step. The obtained results were quite good and easely implemented in real time. In this paper we give emphasis to the optimization of the obtained initial segmentation. As future work we intend to include the optimization step in our video implementation. To do this we need to enhance the performance of the graph-cut optimization. The video implementation is also a good reason to use active light instead of simple background subtraction.

7 Acknowledgments

Anselmo Montenegro is grateful to FAPERJ for the financial support under grant E-26/152.827/2004.

References

[1] A. Agrawala, M. Doncheva, M. Agrawala, S. Drucker, A. Colburn, B. Curless, D. Salesin, and M. Cohen. Interactive digital photomontage. In *Computer Graphics Proceedings ACM SIGGRAPH*, pages 294–302, 2004.

[2] Y. Boykov and M. Jolly. Interactive graph cuts for optimal boundary and region segmentation of objects in n-d images. In *Proceedings of ICCV*, 2001.

[3] Y. Boykov, O. Veksler, and R. Zabih. Markov random fields with efficient approximation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 648–655, 1998.

[4] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on PAMI*, 23:1222–1239, 2001.

[5] L. D. Cohen. On active contour models and ballons. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 53:211–218, 1991.

[6] E. Eisemann and F. Durand. Flash photography enhancement via intrinsic relighting. *Computer Graphics Proceedings ACM SIGGRAPH*, 2004.

[7] S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restorations of images. *IEEE Transactions on Pattern Analysis and Machine Inteligence*, 6:721–741, 1984.

[8] D. Greig, B. Porteous, and A. Seheult. Exact maximum a posteriori estimation for binary images. *J. Royal Statistical Soc.*, pages 271–279, 1989.

[9] H. Ishikawa and D. Geiger. Oclusions, discontinuities, and epipolar lines in stereo. In *Fifth European Conference on Computer Vision, (ECCV'98)*, Freiburg, Germany, 2-6 June 1998.

[10] H. Ishikawa and D. Geiger. Segmentation by grouping junctions. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 125–131, 1998.

[11] M. Kass, A. Witkins, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision and Graphics*, 2:321–331, 1998.

[12] V. Kolmogorov and R. Zabih. Multi-camera scene reconstruction via graph cuts. In *European Conference on Computer Vision*, May 2002.

[13] V. Kolmogorov and R. Zabih. What energy functions can be minimized via graph cuts? *Proc. IEEE Transactions on Pattern Analysis and Machine Inteligence*, 26:147–159, 2004.

[14] V. Kwatra, A. Schdl, I. Essa, G. Turk, and A. Bobick. Graphcut textures: Image and video syntesis using graph cuts. In *ACM Transactions Graphics, Proc. SIGGRAPH*, July 2003.

[15] E. N. Mortensen and W. A. Barrett. Interactive segmentation with intelligent scissors. *Graphical Models and Image Processing*, 60:349–384, 1998.

[16] S. Osher and J. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on hamilton-jacobi formulation. *Graphical Models and Image Processing*, 79:12–49, 1998.

[17] G. Petschnigg, M. Agrawala, H. Hoppe, R. Szeliski, M. Cohen, , and K. Toyama. Digital photography with flash and no-flash image pairs. *Computer Graphics Proceedings ACM SIGGRAPH*, pages 664–672, 2004.

[18] R. Raskar, K. Tan, R. Feris, J. Yu, and M. Turk. Non-photorealistic camera: Depth edge detection and stylized rendering using multi-flash imaging. *Computer Graphics Proceedings ACM SIGGRAPH*, pages 679–688, 2004.

[19] C. Rother, V. Kolmogorov, and A. Blake. Grabcut - interactive foreground extraction using iterated graph cuts. In *Computer Graphics Proceedings ACM SIGGRAPH*, pages 309–314, 2004.

[20] S. Roy and I. Cox. A maximum-flow formulation of the n-camera stereo correspondence problem. In *Int. Conf. On Computer Vision, ICCV'98*, Bombay, India, 1998.

[21] S. Roy and M.-A. Drouin. Non-uniform pyramid stereo for large images. In *IEEE Workshop on Stereo and Multi-Baseline Vision*, Kauai, Hawaii, 2001.

[22] A. Sá, M. Vieira, P. Carvalho, and L. Velho. Range-enhanced active forground extraction. In *Proc. ICIP*, 2005.

[23] J. Shi and J. Malik. Normalized cuts and image segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 731–737, 1997.

[24] A. Smith and J. Blinn. Blue screen matting. *Computer Graphics Proceedings ACM SIGGRAPH*, pages 259–268, 1996.

[25] J. Wang, P. Bhat, R. Colburn, M. Agrawala, , and M. Cohen. Interactive video cutout. *Computer Graphics Proceedings ACM SIGGRAPH*, 2005.