

Isolating Motion and Color in a Motion Blurred Image

Alessandro Giusti Vincenzo Caglioti
Dipartimento di Elettronica e Informazione, Politecnico di Milano
P.za Leonardo da Vinci, 32
20133 Milano - Italy
{giusti,caglioti}@elet.polimi.it

Abstract

Photographic images of moving objects are often characterized by motion blur; analyzing motion blurred images is problematic since the moving object boundaries appear fuzzy and seamlessly blend with the background. In extreme cases, when the object motion is fast in relation to the exposure time, the blurred object image becomes an elongated, semitransparent smear.

We consider a motion-blurred color image of an object moving over a still background: we introduce meaningful entities, the “alpha map” and the “color map”, which bear information about the object motion during the exposure, and its color and texture; we draw connections to the well-known alpha matting problem, providing an original interpretation in this context; we present an analytic technique for extracting the two maps under assumptions on the background and object colors, and explore the relaxation of these assumptions. We provide experimental results on both synthetic and real images, which confirm the correctness of our approach, and describe diverse application examples in fields spanning from 3D reconstruction to image/video enhancement.

1 Introduction

When a moving object is photographed by a still camera, its image is motion blurred because its position changes during the exposure time. If the exposure time is not short enough in relation to the object apparent speed, the object results in a visible “smear” or “streak” in the image, and its contours blend with the background confusing traditional computer vision techniques.

In this paper, the blurred image is analyzed by producing two pieces of information for each image pixel:

- how long the moving object projected there during the exposure time ($\alpha(p)$);
- the color the pixel would have if the background was black ($o(p)$).

Over the whole image, they comprise an “alpha map” and a “color map” (see figure 1). We show that this representation is strictly related to the deeply-studied *alpha matting* problem.

We describe an analytical procedure for computing such maps, in a single blurred color image of a monochromatic object moving over a known background. A monochromatic object is an object whose surface is pigmented with a single color and, possibly, its darker shades – including black; we also extend the technique in order to handle bichromatic objects, pigmented with two different colors, their combinations and darker shades; many common objects, including most types of sport balls can be considered either monochromatic or bichromatic. We show that any further relaxation of the assumptions makes the problem underconstrained: we are considering the most general scenario in which the problem has an analytical solution.

We maintain that, once known, the alpha and color maps can be a valuable aid in the understanding of the blurred object image. Other than applications for image enhancement – for example to highlight images of fast-moving objects (see section 4 and additional material), or to recover the color of a very blurred monochromatic object, the alpha map can be exploited for reconstructing the fast motion of objects which appear blurred, or as a necessary preprocessing step for precisely deblurring a moving object on a sharp background; actually, many works exploiting motion blur of objects as a source of information would benefit of our provided separation of background, object exposure time and object color/texture. In section 5, we briefly explain some of our current work in the field.

The paper is structured as follows: after referencing related works, we formally define the problem of reconstructing the motion blur of a moving object over a known background and highlight analogies to the alpha matting problem (section 2); we provide a theoretically founded, analytical solution to the problem under the assumption that the object is monochromatic, then we explore extensions and limits of the technique (section 3); we then show several practical applications and experimental results (section 4) and finally draw conclusions and outline future developments (section 5).




Figure 1: The image with the (barely visible) motion blurred object (C , at the center) can be interpreted as the temporal average over the exposure time of infinite still images I_t (left). Ordinary background subtraction (image D , after equalization) does not help in characterizing the blur. Our technique provides a map α (analogous to an alpha matte) summarizing the object motion and an object color map o .

1.1 Related Works

Motion blur analysis has been exploited for a number of applications: for example, in [14, 15] quantitative measurements of blur are used in order to estimate the speed of vehi-

cles and spherical objects: the blur parameters are estimated with gradient-based methods, without the support of an exact appearance model of the motion-blurred object over the background. In [6], the curved trajectory of a moving ball is reconstructed from a single, long-exposure image: in this case, an accurate analysis of the blurred ball streak is necessary to find its contours with sufficient precision for the subsequent 3D reconstruction step; this analysis is performed under some very restrictive assumptions, such as constant intensity of the ball image (*i.e.*, no shading), which are not needed when using our proposed technique; very similar considerations hold for [5], which analyzes a slightly-blurred ball and reconstructs its position and velocity. Recently, interesting related applications have also been proposed in [12], where rotational blur estimation is the basis for a visual gyroscope, and in [16], which exploit motion blur to obtain depth information from a single blurred image. Also, in [8] motion blur is used for joint reconstruction of geometry and photometry of scenes with multiple moving objects. In general, these applications take advantage of the additional information that a single motion blurred image incorporates about the scene structure, which can not be extracted from a single still image.

Many techniques ([4, 2], and recently [9]) have been proposed for estimating motion blur parameters, mainly aimed at image “deblurring” (the first attempts date back to 1967 in [20]); spatially-variant motion blur has also been considered in many works, such as [11]. Our work is complementary to these techniques since we do not directly aim at interpreting the motion blur direction or extent; for example, our technique inherently handles complex, nonlinear trajectories or even object deformations; we operate at a lower level, separating the blurred object image from the background, and providing an “alpha map” which isolates the object motion, and a “color map” which is a blurred image of the object, separated from the background. We believe that in some settings, deblurring could actually take advantage of our resulting representation (see section 5).

We show in section 2.1 that part of our problem can be reduced to *alpha matting* (or *layer extraction*): a classic problem in computer vision which has been shown to be unconstrained in the general case; the subject has been extensively studied due to its immediate, obvious applications in many fields, and various solutions have been proposed: some ([21, 17]) require a specific background (blue screen matting), whereas others ([3, 7, 19, 22, 13, 1]), with minimal user assistance, handle unknown backgrounds (*natural image matting*) by means of sophisticated heuristics. None of these algorithms was designed for motion-blurred images – for example, all require that a part of the image has $\alpha = 1$, and many assume that mixed pixels are rare; nonetheless, some actually work acceptably in specific instances of this unforeseen scenario (see additional material), but quantitative use of the resulting alpha map is rarely an option.

2 Definitions and model

A motion blurred image is obtained when the scene projection on the image plane changes during the camera exposure period $e = [t_0, t_0 + \Delta t]$. The final image C is obtained as the integration of infinite sharp images, each exposed for an infinitesimal portion of e . An equivalent interpretation, which is more meaningful in this setting, is considering the motion blurred image as the temporal average of infinite sharp images I_t , each taken with the same exposure time Δt but representing the scene frozen at a different instant $t \in e$

(see figure 1). This technique is also implemented in many 3D rendering packages for accurate synthesis of motion blurred images.

If the camera is static and a single object is moving in the scene, the static background in the final image is sharp since its pixels are of constant intensity in each I_t ; conversely, the image pixels which are affected by the moving object, possibly only in some of the I_t images, belong to the motion-blurred image of the object.

2.1 Motion blur and alpha compositing

For a pixel p , define $i(p) \subseteq e$ the time interval in e during which p belongs to the object image. We define $\alpha(p)$ as the fraction of e during which the object projects to p :

$$\alpha(p) = \|i(p)\|/\Delta t. \quad (1)$$

Let $B(p)$ be the intensity of the background at p . Since C is the temporal average of all the I_t images, $C(p) = \alpha(p)o(p) + (1 - \alpha(p))B(p)$, where $o(p)$ is the temporal average over $i(p)$ of the intensity of image pixel $C(p)$:

$$o(p) = \frac{1}{\|i(p)\|} \int_{t \in i(p)} I_t(p) dt. \quad (2)$$

To sum up, the intensity of a pixel p in the motion blurred image C can be interpreted as the convex linear combination of two factors: the "object" intensity $o(p)$, weighted $\alpha(p)$, and the background intensity. The resulting equation is the usual *over* Porter-Duff alpha compositing equation [18] for a pixel with transparency $\alpha(p)$ and intensity $o(p)$ over the background pixel $B(p)$.

$o(p)$ can be interpreted as the intensity that p would have in the motion blurred image over a black background, rescaled by a $\frac{1}{\|i(p)\|}$ factor. $o(p)$ is meaningless if p is not affected by the object image during e (i.e., if $\alpha(p) = 0$).

The considerations expressed so far can be directly applied to a color image, considering all channels separately, but noting that $\alpha(p)$ is constant along all the channels:

$$\begin{aligned} C_r(p) &= \alpha(p)o_r(p) + (1 - \alpha(p))B_r(p) \\ C_g(p) &= \alpha(p)o_g(p) + (1 - \alpha(p))B_g(p) \\ C_b(p) &= \alpha(p)o_b(p) + (1 - \alpha(p))B_b(p). \end{aligned} \quad (3)$$

2.2 Assumptions on the object color

In order to provide an analytic reconstruction technique, we require that the object surface is monochromatic, meaning that, in the HSV color space, its hue and saturation are fixed – whereas different brightness values (shades) of the color are allowed. In the RGB color space, this translates to requiring that the object surface colors all lie on a single line passing through $(R, G, B) = (0, 0, 0)$. In practice, this includes many real-world objects; note that black surface parts of the object are always allowed.

Under this assumption and ignoring possible mixed lighting, specularity or transparency of the object surface, the colors of the object image satisfy the same condition even if shading is taken into account: they all still lie on a single line passing through $(R, G, B) = (0, 0, 0)$ – whereas their actual hue and saturation depend on the white balance and light color. We call this line in RGB space l .

In the motion-blurred image, $o_{rgb}(p)$ is constrained to lie on l as well, because it is a linear combination of colors belonging to l ; if l is defined in polar coordinates:

$$\begin{aligned} o_r(p) &= k \cdot \sin \phi \cos \theta \\ o_g(p) &= k \cdot \cos \phi \sin \theta \\ o_b(p) &= k \cdot \cos \phi. \end{aligned} \quad (4)$$

3 Analysis technique

3.1 Recovering the object color




Figure 2: l (red line) passes through the origin and its orientation can be retrieved from the colors of C and B in two points. $o(p)$ is found as the intersection between l and the line passing through $C(p)$ and $B(p)$.

In a motion-blurred image, the object hue and saturation can be easily recovered if at least one pixel p exists such that $\alpha(p) = 1$; in other words, if p belongs to the image of the object during the whole exposure period, and is therefore not affected by the background color. In this case, the hue and saturation of $C(p)$ univocally define l in the RGB space.

Else, no pixel in the blurred image has a color which can be directly related to l . Therefore, we consider how the color of pixels belonging to the motion blurred images change between the background image B and the motion-blurred image C . Given two pixels p_1 and p_2 belonging to the object streak, we can write a system in 12 equations and 12 unknowns ($k_1, k_2, \phi, \theta, \alpha(p_1), \alpha(p_2)$, the 6 $o_{rgb}(p_1)$ and $o_{rgb}(p_2)$ values), which has a single solution (see Appendix in supplementary material). A much simpler interpretation is given below.

Geometric interpretation in RGB space For each pixel p , from (4) we know that $o(p)$ lies on the line in RGB space connecting $C(p)$ to $B(p)$. However, due to the object non-uniform color and shading, $o(p)$ depends on the actual pixel. Therefore lines in RGB space defined by $C(p_1), B(p_1)$ and $C(p_2), B(p_2)$ will, in general, not intersect.

We know from (2) that $o(p_1)$ and $o(p_2)$ both lie on the same line in RGB space, which also passes through black ($(R, G, B) = (0, 0, 0)$). Therefore, we compute l as the only line passing through black and intersecting lines defined by $C(p_1)B(p_1)$ and $C(p_2)B(p_2)$; in other words, l can be found as the intersection of two planes, defined by RGB points $C(p_1), B(p_1), (0, 0, 0)$ and $C(p_2), B(p_2), (0, 0, 0)$. This requires that $B(p_1)$ and $B(p_2)$ are

linearly independent vectors in RGB space (*i.e.*, they do not represent the same color with different shading).

3.2 Recovering object exposure time and average color for each pixel

For a pixel in the blurred image, we can compute $\alpha(p)$ and $o(p)$ if l is known:

$$\begin{aligned}
 o_r(p) &= k \cdot \sin \phi \cos \theta \\
 o_g(p) &= k \cdot \cos \phi \sin \theta \\
 o_b(p) &= k \cdot \cos \phi \\
 C_r(p) &= \alpha(p)o_r(p) + (1 - \alpha(p))B_r(p) \\
 C_g(p) &= \alpha(p)o_g(p) + (1 - \alpha(p))B_g(p) \\
 C_b(p) &= \alpha(p)o_b(p) + (1 - \alpha(p))B_b(p).
 \end{aligned} \tag{5}$$

It's 6 equations for 5 unknowns ($\alpha(p)$, k and the 3 $o_{rgb}(p)$ values), so the system is overdetermined.

Geometric interpretation in RGB space From equation (3), $C(p)$ is a linear combination of $B(p)$ and $o(p)$: therefore, the line in RGB space connecting $B(p)$ and $C(p)$ intersects l in $o(p)$. This identifies $o(p)$. Once $o(p)$ is known, $\alpha(p)$ is found as $\frac{B(p)C(p)}{B(p)o(p)}$.

Pixels belonging to the background are unchanged between $B(p)$ and $C(p)$, therefore $\alpha = 0$ unless $B(p)$ belongs to l , which is a degenerate case. Contrarily, if a pixel p belongs to the object image during the entire exposure time, $C(p)$ belongs to l , so $\alpha = 1$.

3.3 Extensions and limits

Computing $o(p)$ and $\alpha(p)$ using (5) is an overconstrained problem. This suggests that the same theoretical framework could be adapted to handle more general problems and/or relax the assumptions.

- The object monochromaticity assumption can be relaxed: the problem can be still solved if the object surface is bichromatic – meaning that it is composed of two different colors, their shades towards black, and their (convex) linear combinations. Therefore, the surface colors of the object lie on a plane in RGB space. Since that plane passes through black, also the possible colors of the shaded object in a sharp image belong to a plane π (albeit possibly different, if light is not white). In the blurred image, $o(p)$ colors, obtained as linear combination of colors belonging to π , still lie on π .
- The nonspecular shading assumption can be relaxed: specular highlights appearing on a monochromatic object surface would cause the space of possible $o(p)$ values to extend on a plane, similarly to the previous case.
- The uniform lighting assumption can be relaxed: if only two different colors of light sources are present in the scene, all $o(p)$ colors still lie on a plane.

In practice, any of these three ways to exploit the additional degree of freedom make the reconstruction procedure for α an exactly determined problem, instead of overdetermined. Algebraically, one of the equations constraining the $o(p)$ position is removed; in the geometric interpretation, $o(p)$ is found as the intersection between a line (passing through $C(p)$ and $B(p)$) and a plane (π) – instead of the intersection between two lines. The degenerate case is when $B(p)$ belongs to π .

However, unless some additional assumptions or approximations are introduced, there is no procedure for automatically determining π analogous to the one introduced in section 3.1 for finding l . In practical applications, π can be computed in a different, sharp image of the object, by directly picking two colors linearly independent in RGB; they could be also picked directly in the blurred image, if some pixels exist for which $\alpha(p) = 1$. Using more than two pixels would also allow a more robust estimation of π by least-squares or RANSAC.

Unfortunately, the additional degree of freedom can not be exploited to handle variations in background lighting, such as shadows. In fact, this can be interpreted geometrically as constraining that $B(p)$ lies on a line in RGB passing through $(0,0,0)$. This line lies on the same plane defined by l and $C(p)$; this explains why the resulting system of equations is not determined. However, if the background color $B(p)$ is allowed to change along a line which does not pass through $(0,0,0)$, $o(p)$ and α can still be computed. This has a practical application for pixels near a sharp discontinuity between two different background colors: these are often problematic because of several factors, including image compression and minimal camera rotations after the background has been sampled. By only constraining $B(p)$ to lie on a line passing through both background colors, this problem can be avoided.

So far, we assumed that the object in each of the I_t images is perfectly sharp, and we ignored the presence of “mixed pixels”, originating from imperfect focus, and anyway always present along the object contours – a well known problem in matting research. However, our procedure inherently handles these cases, and works with defocused, “fuzzy” or even semitransparent objects¹.

4 Experimental results and application examples




Figure 3: Synthetic example. From left to right: still bichromatic object, motion blurred image, reconstructed alpha map, reconstructed color map. Object composited over a different background using computed alpha and color map.

The technique has been evaluated both with synthetic images and with real images.

In figure 3, a bichromatic motion blurred cube is analyzed. Pixels in the central part of the streak belong to the object image during the whole exposure time, and are used

¹Object transparency is handled under rather restrictive assumptions such as absence of refractive effects

to determine the plane of the possible object colors. The absence of any error in the reconstructed alpha and color maps confirms the theoretical validity of our approach. Once the alpha and color maps are known, the blurred object can be composited over a different background as shown; other interesting applications include artificial opacity or color manipulation – *e.g.* for visualizing long trajectories of balls in long-exposure photos.




Figure 4: *First four lines.* From left to right: original image, alpha map, color map. Last object is bichromatic teddy bear, others are monochromatic small sheets of colored paper. Note that the color map is meaningless outside the object blurred image. *Last line:* part of the trajectory of a moving table tennis ball; from left to right: original image, alpha map, color map, enhanced image obtained by compositing a linearly-transformed color map over the background, using a multiplied alpha map for enhancing visibility (note preserved shading). See additional material for other examples and applications.

Figure 4 shows examples with real images. The reconstruction of a monochromatic object's color works reliably and robustly: in tests with table tennis balls and small colored paper sheets we were able to retrieve the color of the object image from a long streak over a multicolored background within an Euclidean distance of 3 units in the RGB color cube (each channel having $0 \div 255$ range) in all tests. Reconstructed alpha and color maps are heavily affected by image noise, which is expected since we work one pixel at a time

and do not yet perform any filtering. Also, discrepancies between the sampled and actual background introduce severe artifacts, which can be reduced, at least in part, with the technique described in section 3.3.

The theoretical background we presented builds upon the linearity of color channels, which is not normally preserved in images directly coming from digital cameras, since nonlinear gamma curves are applied automatically by the camera firmware unless the image is shot in RAW format. More generally, doing quantitative processing on the image colors requires that the camera response function is known. The estimation of the camera response function is the subject of several studies such as [10], which provide various convenient solutions which allow to linearize the intensity values of the color channels.

Often the reconstructed object color $o(p)$ exceeds the dynamic range of the image. This is fine, and means that, if the ball image projected to that pixel for the whole exposure time, it would have been overexposed; in this regard, note that accurate rendering of motion blur is an important application of High Dynamic Range image synthesis.

5 Conclusions and Future Developments

We formally defined the problem of reconstructing the motion blur of a moving object over a known background, describing how it relates to alpha matting, a deeply explored topic in computer vision.

If the object is monochromatic, we can retrieve its color even if it only appears as a semitransparent smear. If the object is monochromatic or bichromatic and its color is known, we can compute the fraction of the exposure time during which the object projection overlapped each pixel; also, we can determine how the image would appear if the background was black.

This allows, for example, to composite the blurred object over any background, artificially change its opacity or color or study patterns left by differently pigmented parts. Other applications which could take advantage from these techniques are related to the reconstruction of the object's motion from a single blurred image, and point spread function estimation problems for image restoration and deblurring.

We are currently applying these techniques in practice, while developing algorithms to robustly exploit them in presence of noise and other nonidealities. We are also using these techniques for motion estimation and 3D reconstruction from a single motion blurred image, by exploiting the alpha map's interpretation as the fraction of exposure time during which the object image overlapped that pixel. This allows, among others, applications in temporal superresolution of apparent contours motion. We are also applying known blind deblurring techniques on the alpha and color maps separately, in order to precisely deblur objects on a fixed, sharp background.

References

- [1] N. E. Apostoloff and A. W. Fitzgibbon. Bayesian video matting using learnt image priors. In *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2004.
- [2] B. Bascle, Andrew Blake, and Andrew Zisserman. Motion deblurring and super-resolution from an image sequence. In *Proc. of European Conference on Computer Vision (ECCV)*, pages 573–582, 1996.

- [3] A. Berman, P. Vlahos, and A. Dadourian. Comprehensive method for removing from an image the background surrounding a selected object. U.S. Patent 6,134,345, 2000.
- [4] M. Bertero and P. Boccacci. *Introduction to Inverse Problems in Imaging*. Inst. of Physics Publishing, 1998.
- [5] Giacomo Boracchi, Vincenzo Caglioti, and Alessandro Giusti. Ball position and motion reconstruction from blur in a single perspective image. In *Proc. of International Conference on Image Analysis and Processing (ICIAP)*, 2007.
- [6] V. Caglioti and A. Giusti. Ball trajectory reconstruction from a single long-exposure perspective image. In *Proc. of Workshop on Computer Vision Based Analysis in Sport Environments (CVBASE)*, 2006.
- [7] Y. Chuang, B. Curless, D. Salesin, and R. Szeliski. A bayesian approach to digital matting. In *Proc. of CVPR 2001*.
- [8] Paolo Favaro and Stefano Soatto. A variational approach to scene reconstruction and image segmentation from motion-blur cues. In *Proc. of Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 01, pages 631–637, Los Alamitos, CA, USA, 2004. IEEE Computer Society.
- [9] Rob Fergus, Barun Singh, Aaron Hertzmann, Sam T. Roweis, and William T. Freeman. Removing camera shake from a single photograph. In *ACM SIGGRAPH 2006 Papers*, pages 787–794, New York, NY, USA, 2006. ACM Press.
- [10] Michael D. Grossberg and Shree K. Nayar. Modeling the space of camera response functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(10):1272–1282, 2004.
- [11] S.K. Kang, Y.C. Choung, and J.K. Paik. Segmentation-based image restoration for multiple moving objects with different motions. In *Proc. of International Conference on Image Processing (ICIP)*, pages I:376–380, 1999.
- [12] G. Klein and T. Drummond. A single-frame visual gyroscope. In *Proc. of British Machine Vision Conference (BMVC)*, 2005.
- [13] Anat Levin, Dani Lischinski, and Yair Weiss. A closed form solution to natural image matting. In *Proc. of CVPR 2006*, pages 61–68, Washington, DC, USA, 2006. IEEE Computer Society.
- [14] Huei-Yung Lin. Vehicle speed detection and identification from a single motion blurred image. In *Proc. of IEEE Workshop on Applications of Computer Vision*, pages 461–467, 2005.
- [15] Huei-Yung Lin and Chia-Hong Chang. Automatic speed measurements of spherical objects using an off-the-shelf digital camera. In *Proc. of IEEE International Conference on Mechatronics*, pages 66–71, 2005.
- [16] Huei-Yung Lin and Chia-Hong Chang. Depth recovery from motion blurred images. In *Proc. of International Conference on Pattern Recognition (ICPR)*, volume 1, pages 135–138, Los Alamitos, CA, USA, 2006. IEEE Computer Society.
- [17] Y. Mishima. Soft edge chroma-key generation based upon hexoctahedral color space. U.S. Patent 5,355,174, 1993.
- [18] T. Porter and T. Duff. Compositing digital images. *Computer Graphics*, 1984.
- [19] M. Ruzon and C. Tomasi. Alpha estimation in natural images. In *Proc. of CVPR 2000*.
- [20] D. Slepian. Restoration of photographs blurred by image motion. *Bell System Tech.*, 46(10):2353–2362, December 1967.
- [21] Alvy Ray Smith and James F. Blinn. Blue screen matting. In *SIGGRAPH '96: Proc. of the 23rd annual conference on Computer graphics and interactive techniques*, pages 259–268, New York, NY, USA, 1996. ACM Press.
- [22] Jian Sun, Jiaya Jia, Chi-Keung Tang, and Heung-Yeung Shum. Poisson matting. In *ACM SIGGRAPH 2004 Papers*, pages 315–321, New York, NY, USA, 2004. ACM Press.