

On the Apparent Transparency of a Motion Blurred Object

Alessandro Giusti, Vincenzo Caglioti

► **To cite this version:**

Alessandro Giusti, Vincenzo Caglioti. On the Apparent Transparency of a Motion Blurred Object. Proceedings of the First International Workshop on Photometric Analysis For Computer Vision - PACV 2007, Oct 2007, Rio de Janeiro, Brazil. 8 p. inria-00264748

HAL Id: inria-00264748

<https://hal.inria.fr/inria-00264748>

Submitted on 17 Mar 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On the Apparent Transparency of a Motion Blurred Object

Alessandro Giusti

Vincenzo Caglioti

Dipartimento di Elettronica e Informazione, Politecnico di Milano

P.za Leonardo da Vinci, 32 20133 Milano – Italy

alessandro.giusti@polimi.it vincenzo.caglioti@polimi.it

Abstract

An object which moves during the exposure time results in a blurred smear in the image. We consider the smear as if it was the image of a semitransparent object, and we retrieve its alpha matte by means of known techniques. Our work is focused on the qualitative and quantitative analysis of the alpha matte: in particular, we show that the alpha value at a pixel represents the fraction of the exposure time during which the object image overlapped that pixel; then, we prove that the alpha matte enjoys a number of properties, which can be used to derive constraints on the object’s apparent contour, and its motion, from a single motion-blurred image. We show that some of these properties also hold on the original image. We point out possible applications, and experimental results both on synthetic and real finite-exposure images.

1. Introduction

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

In this paper, we provide a meaningful interpretation of the smear’s alpha matte (i.e., transparency), which can be recovered by means of controlled shooting conditions or known matting algorithms. Then, we prove that the alpha matte of the blurred object image enjoys several properties, which are linked to the object’s apparent contour and its motion during the exposure time.

Other than generically supporting the understanding of the object smear, this allows us to retrieve:

- the apparent contour at the beginning and end of the exposure;

- parts of the contour at arbitrary time instants inside the exposure time;
- envelopes of the moving contour;
- the path of corners on the contour;
- speed discontinuities of the contour, caused for example by impacts of the moving object.

Most literature dealing with motion blur is concerned with removing it (“deblurring”), either assuming that it affects the whole image because of camera shake [5, 19], or aiming at deblurring moving objects [9]. In [8], the study of the blurred object’s alpha matte is used for recovering the point spread function, improving deblurring performance.

In our setting, on the contrary, we want to take advantage of the additional information that a single motion blurred image incorporates about the scene structure and evolution, which can not be extracted from a single still image. Several works in literature follow a similar approach: [12, 14] estimate the speed of a moving vehicle and of a moving ball, respectively; in [10], rotational blur estimation is the basis for a visual gyroscope, and [13] exploits motion blur to obtain depth information from a single blurred image. In [4] motion blur is used for joint reconstruction of geometry and photometry of scenes with multiple moving objects.

Instead of referring to a specific application, our work has a general validity when moving objects are imaged from a still camera, and provides a sound theoretical foundation for interpreting and extracting information from blurred smears; contrarily to most works in the field, we use a general blur model, and do not require that blur is representable in the frequency domain. This allows us to handle any type of motion.

The paper is structured as follows: at the beginning we recall a general model for motion-blurred image formation (section 2), and show how the smear’s transparency can be meaningfully interpreted in this context; we briefly show how transparency can be recovered under varying constraints, then introduce our notation and basic properties. In section 3 we introduce and prove our main results,

which relate features on the alpha matte to meaningful constraints on the moving object’s apparent contours during the exposure time. We discuss practical applications of our theory in section 4, whereas Section 5 presents experimental results on both real and synthetic images. Finally, we provide conclusions and list future works (section 6).

2. Model and definitions

We begin by presenting a model for the image of a moving object, highlighting the connections to the matting problem by introducing the $\alpha(p)$ quantity, which will be our main object of study in the following.

2.1. Motion blur model and its relation with alpha matting

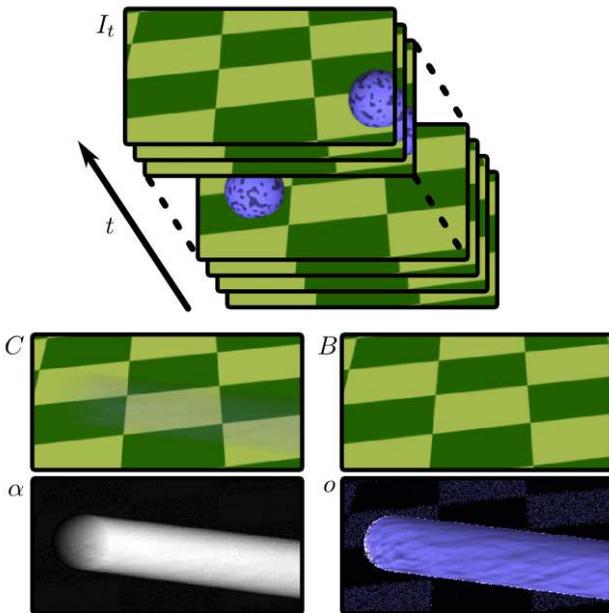


Figure 1. The image C of the motion blurred object can be interpreted as the temporal average over the exposure time of infinite still images I_t (top). We show that the blurred smear can be interpreted as a semitransparent layer, whose alpha matte we analyze in the following.

A motion blurred image is obtained when the scene projection on the image plane changes during the camera exposure period $[t', t'']$. The final image C is obtained as the integration of infinite sharp images, each exposed for an infinitesimal portion of $[t', t'']$. In an equivalent interpretation (figure 1), we can consider the motion blurred image as the temporal average of infinite sharp images I_t , each taken with the same exposure time $t'' - t'$ but representing the scene frozen at a different instant $t \in [t', t'']$. 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 (smear) of the object.

For a pixel p , define $i(p) \subseteq [t', t'']$ the set of time instants during which p belongs to the object image. We finally define $\alpha(p)$ as the fraction of $[t', t'']$ during which the object projects to p :

$$\alpha(p) = ||i(p)|| / (t'' - t'). \quad (1)$$

Let $B(p)$ be the intensity of the background at p . Since C is the temporal average of all the I_t images, $C(p) = \alpha(p)o(p) + (1 - \alpha(p))B(p)$. $o(p)$ 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 well-known Porter-Duff alpha compositing equation [16] for a pixel with transparency $\alpha(p)$ and intensity $o(p)$ over the background pixel $B(p)$. Recovering $\alpha(p)$ from a motion-blurred image is therefore equivalent to the well-known *layer extraction* (or *alpha matting*) problem.

2.2. Recovery of the alpha matte

If the object appears with a given constant intensity on a known background, $\alpha(p)$ can be readily derived at each pixel by solving a first-degree linear equation.

In a color image, we can extend our considerations to all channels separately, noting that $\alpha(p)$ is constant along all the channels:

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

Then, if the background is known and the object colors meet some constraints, $\alpha(p)$ can analytically be computed in the whole image [7].

In the general case, the matting problem is underconstrained, even if the background is known: still, in literature many algorithms have been proposed : some ([20, 15]) require a specific background (blue screen matting), whereas others ([2, 3, 17, 21, 11, 1]), with minimal user assistance, handle unknown backgrounds (natural image matting) and large zones of mixed pixels ($0 < \alpha < 1$). Although none is explicitly designed for the interpretation of blurred smears, we can get satisfactory results in our peculiar setting. As we

show in this paper, the alpha matte of blurred smears enjoys some rather restrictive properties, and several constraints can also be enforced for the object color: these additional constraints could be integrated in existing algorithms in the future, for a more accurate extraction of the alpha matte of blurred smears.

As we discuss in section 4.2, some of our results also hold for the original blurred image, which allows applications even without matting.

2.3. Definitions, basic assumptions and properties

During the exposure time $[t', t'']$, the object contour changes. Let $c(t)$ be the object apparent contour at t .

We partition the image in regions, which we classify on the basis of two criteria:

- Whether the region is inside $c(t')$ (R^\bullet) or not (R°), which we indicate with a superscript.
- How many times each image point belonging to the region is crossed by c during the exposure time, which we indicate with a subscript.

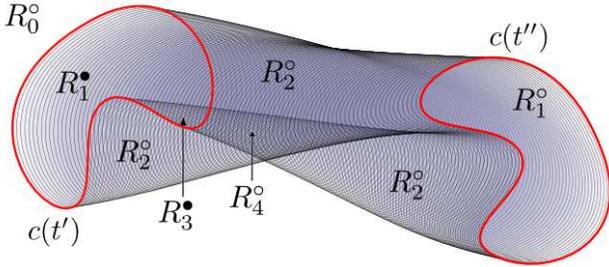


Figure 2. The smear of a rotating object and its regions.

The background ($\alpha(p) = 0$) is the R_0° region. On the contrary, the R_0^\bullet region is composed by all and only the pixels covered by the object during the whole exposure ($\alpha = 1$); in figure 2, there is no R_0^\bullet region.

Note that if any point inside $c(t')$ is crossed once by the apparent contour, it will be outside $c(t'')$, and vice-versa; this is easily generalized to the following property.

Property 1. An R_{2n+1}^\bullet or R_{2n}° region is outside $c(t'')$. An R_{2n}^\bullet or R_{2n+1}° region is inside $c(t'')$.

3. Properties of the blurred smear alpha matte

We present a meaningful property of α values in R_1 regions, then discuss region topology and classify borders between regions. Finally, we show how this relates to the alpha matte.

3.1. α inside R_1 regions

Theorem 2. The iso- α curve $\alpha = \bar{\alpha}$ within an R_1^\bullet region is part of the apparent contour $c(t' + \bar{\alpha}(t'' - t'))$. Similarly, the iso- α curve $\alpha = \bar{\alpha}$ within an R_1° region is part of $c(t'' - \bar{\alpha}(t'' - t'))$.

Proof. Each pixel p belongs to the object image for exactly $\alpha(p)(t'' - t')$ seconds. In an R_1^\bullet region each pixel is crossed once by the object contour, therefore it belongs to the object image during a compact time interval starting at t' and ending when it is crossed by the object contour, at time $t' + \alpha(p)(t'' - t')$. The first part of the thesis immediately follows. A similar argument holds for the second part of the thesis. \square

Corollary. In an R_1 region the direction of $\nabla\alpha(p)$ is orthogonal to the object contour in p at time $\alpha(p)(t'' - t')$; $|\nabla\alpha(p)|$ is inversely proportional to the apparent contour motion speed component along the direction of $\nabla\alpha(p)$. In an R_1° region $\nabla\alpha(p)$ points towards the motion direction, whereas in R_1^\bullet , $\nabla\alpha(p)$ points towards the opposite direction.

3.2. Space-time representation of $c(t)$

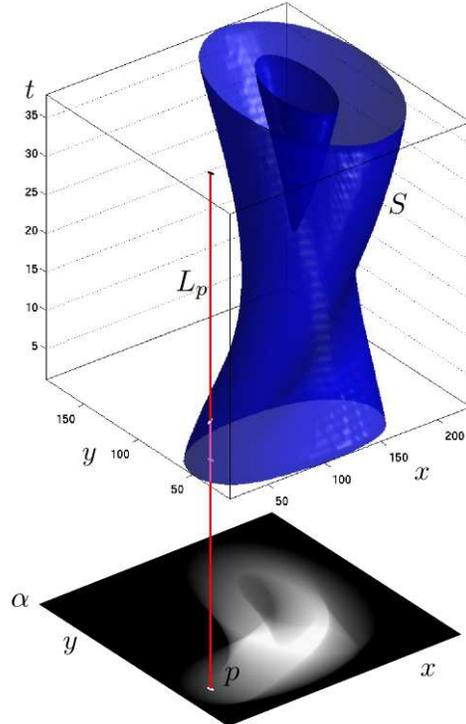


Figure 3. Line L_p intersects S once, since p belongs to an R_1 region.

Consider a 3D solid K in the (x, y, t) space, such that each cross section of K at $t = \bar{t}$ coincides with the image re-

gion inside $c(\bar{t})$. K is bounded by $t = t'$ and $t = t''$ planes, and summarizes the evolution of the object contours during the exposure¹. We also refer to S (see figure 3), which is the lateral surface of K , i.e. the surface of K except the upper ($t = t'$) and lower ($t = t''$) caps²; the intersection of S with a $t = \bar{t}$ plane is $c(\bar{t})$.

Consider a point $p = (p_x, p_y)$ on the x, y plane, and a line L_p in the x, y, t space, defined as $x = p_x, y = p_y$. The number of intersections of S with L_p coincides with the number of times image point p is crossed by c during the exposure time.

Definition 3. An image point p is a degenerate point if and only if none of the following conditions hold:

- (1) L_p is not tangent to S , and p does not belong to either $c(t')$ or $c(t'')$.
- (2) L_p is tangent to S at a single point, and p does not belong to either $c(t')$ or $c(t'')$.
- (3) L_p is not tangent to S , and p belongs to $c(t')$ but not to $c(t'')$.
- (4) L_p is not tangent to S , and p belongs to $c(t'')$ but not to $c(t')$.

The physical meaning of said degeneracy is described in the remarks to the following theorem.

3.3. Borders between regions

Definition 4. We classify borders between regions as follows:

type 1a: connects an R_n region with an R_{n+1} region, inverting the superscript;

type 1b: connects an R_n region with an R_{n+1} region, preserving the superscript;

type 2: connects an R_n region with an R_{n+2} region, always preserving the superscript.

Theorem 5. *If apparent contours move with finite speed and a discrete set of degenerate image points exists:*

- i *All possible borders between regions can be classified according to definition 4; multiple borders may overlap at a degenerate point.*
- ii *The union set of all type 1a borders coincides with $c(t')$; the union set of all type 1b borders coincides with $c(t'')$.*

¹A similar representation has been previously proposed for action classification in [22]

² S can be composed by several isolated parts (as in the figure), but we assume that each is smooth everywhere except possibly in a 0-measure set (ridge)

iii *A type 2 border is the envelope³ of $c(t)|_{t \in [t_1, t_2] \subseteq [t', t'']}$;*

Proof. Consider a non-degenerate image point. Then, only one of the conditions in definition 3 holds.

- If condition (1) holds, then p is not part of a region border. In fact, all points p' in p 's neighborhood belong to the same region as p , because their $L_{p'}$ lines in x, y, t space are crossed by S the same number of times; points p' they are all either inside or outside $c(t')$ – thus the region superscript is kept.
- If condition (2) holds, p belongs to a type 2 border. In fact, in p 's neighborhood two points p' and p'' exist such that $L_{p'}$ crosses S k times, and $L_{p''}$ crosses S $k + 2$ times. The region superscript is kept because all points in the neighborhood are either inside or outside $c(t')$.
- If condition (3) holds, p belongs to a type 1a border. In fact, in p 's neighborhood two points p' and p'' exist such that $L_{p'}$ crosses S k times, and $L_{p''}$ crosses S $k + 1$ times. p' and p'' are at opposite sides of $c(t')$, thus the region superscript is inverted.
- If condition (4) holds, p belongs to a type 1b border. p' and p'' are not at opposite sides of $c(t')$, thus the region superscript is preserved.

□

Remark. Degenerate points are points belonging to more than a single border; by assuming a discrete set of degenerate points, we exclude that two borders coincide along a finite curve. In practice, this condition is nearly always met. Examples of violating scenes are an object moving with a trajectory which exactly retraces itself, or an object with a straight horizontal contour part translating horizontally (a segment of a type 1a border overlaps a part of a type 2 border). However, even in scenes not meeting the constraints,

3.4. Interpretation of features on the alpha matte

Features on the alpha matte can be related to region borders by means of the following theorem.

Theorem 6. *Under the same broad assumptions as theorem 5, and additionally excluding the following degenerate cases:*

- *presence of cusps on S ;*
- *presence of a ridge on S whose projection on the x, y plane overlaps along a finite curve with a region border or with the projection of a different ridge;*

³If $c(t)$ is not smooth, (part of) a type 2 border may represent the locus of points touched by a corner translating along a direction not included by the corner itself. This slightly broadens the definition of envelope.

let $\nabla\alpha$ be the gradient of α ; if p belongs to a single region border then $\nabla\alpha$ is discontinuous in p .

Proof sketch. The theorem can be proven exploiting the space-time representation of $c(t)$ introduced in section 3.2. $\alpha(p)$ can be interpreted as the total length of P 's intersection with K . When p crosses a border of type 1, L_p crosses a ridge on the surface of K , which reflects to a discontinuity in $\nabla\alpha$: our hypotheses exclude that the discontinuity may be cancelled by another ridge in K . When p crosses a border of type 2, P is tangent to K 's surface, which reflects to a discontinuity in $\nabla\alpha$ as well, unless S has a cusp in the tangency point, which is excluded by our hypotheses. \square

Remark. In practice, the degenerate cases in the theorem statement almost never occur; they exclude that a discontinuity in $\nabla\alpha$ is neutralized by a corresponding opposite source of discontinuity.

The reverse of theorem 6 also holds under the following much more restrictive assumptions, which are met by the perspective image of a strictly convex smooth solid in free motion:

1. smooth, strictly convex apparent contours;
2. apparent contours moving with continuity in speed.

Then, $\nabla\alpha$ has a jump discontinuity always and only across a region border. In particular, at a type 1 border we have a jump discontinuity of $\nabla\alpha$ in which both the left and right limits have finite norm. At a type 2 border, one of the left or right limits of $\nabla\alpha$ has infinite norm.

From the proof of theorem 6 we easily derive that also any abrupt change in the speed or orientation of (a part of) apparent contour results in a discontinuity in $\nabla\alpha$. This occurs:

- when the object suddenly changes its motion at a time \bar{t} , e.g. as the result of an impact; in this case a discontinuity in $\nabla\alpha$ overlapping with $c(\bar{t})$ is originated.
- when the moving object is not strictly convex: then a discontinuity in $\nabla\alpha$ at point p occurs as p 's viewing ray is bitangent to the object at a given $t \in [t', t'']$.
- along the path of a corner on the apparent contour.

4. Discussion and practical considerations

4.1. Classifying borders and regions

Theorem 6 allows to locate a superset of all region borders by looking for discontinuities in $\nabla\alpha$. Type 2 borders, when representing the envelope of smooth contours, are easily classified because $\nabla\alpha$ has an extremely large norm just besides the border. If the restrictive assumptions hold, which ensure the validity of the reverse of theorem 6, then

there is a one-to-one relation between discontinuities and borders. Else, some discontinuities may be spurious, and e.g. may represent the path of corners in the apparent contour which do not originate a legitimate region border, or outline the contour when the object abruptly changes its speed.

R_0° and R_0^\bullet regions can be initialized where $\alpha = 0$ and $\alpha = 1$, respectively.

The type 1 borders constitute the union of $c(t')$ and $c(t'')$. Assuming that the two can be separated, there is an uneliminable ambiguity about which is which, so one is arbitrarily chosen as $c(t')$ and classified as Type 1a; all contained regions are set as R^\bullet ; all remaining regions are R° . As a result of such uneliminable ambiguity, *two consistent region labelings are always possible, corresponding to opposite motion directions of the object*; one can be transformed into the other by simply inverting the superscript of odd regions.

The region subscripts are not affected by the ambiguity, and can be derived by means of the constraints stated by theorem 5; we verified on synthetic images that in most images only one consistent labeling (plus its ‘‘inverted’’ twin) satisfies all such constraints.

4.2. Operating on the original image

Under broad assumptions, discontinuities in $\nabla\alpha$ result in discontinuities in the gradient of the original image. Locating region borders and other features directly from the image gradient is therefore theoretically possible. Practical feasibility may be hindered by the possible interference of the background and object shading and texture, and greatly depends on the application scenario.

Figure 4 and Figure 8 illustrate examples of this possibility on real images.

4.3. Applications

By locating and classifying regions, their borders, and other discontinuities of $\nabla\alpha$, strong constraints on the object apparent contours during the exposure can be reconstructed. This allows to extract a wealth of motion information from the object smear, which would not be available in a sharp image.

Independent of the specific application scenario, theorem 2 allows to partially reconstruct apparent contours at intermediate time instants belonging to the exposure period, as the union set of an iso- α curve with value $\bar{\alpha}$ within R_1^\bullet and an iso- α curve with value $1 - \bar{\alpha}$ within R_1° . This can also be applied for temporal superresolution of apparent contours in a video whose frames are exposed for a significant part of the inter-frame interval; this has analogies with Shechtman and Caspi's work [18], which tackles a similar problem with a completely different approach. Another

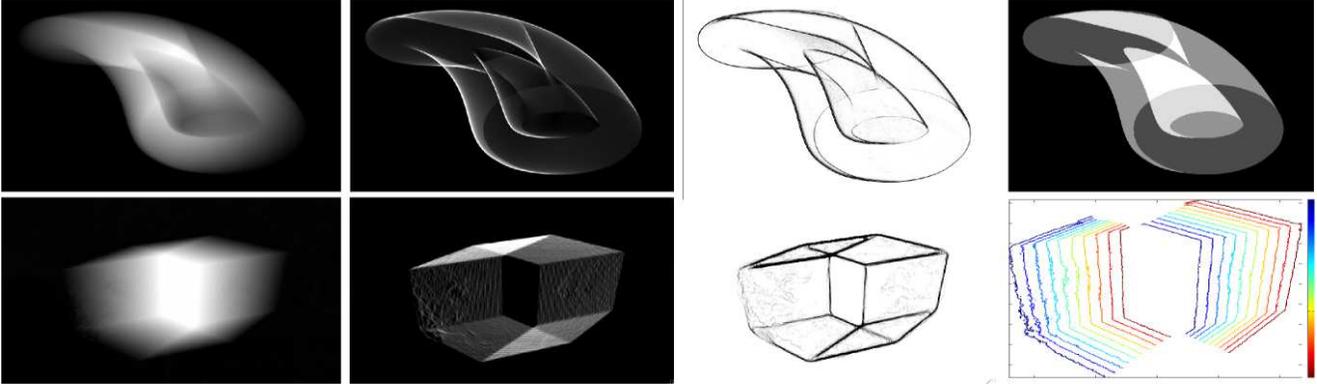


Figure 5. Results on synthetic images. *First row*: rototraslating torus; from left to right: α ; $|\nabla\alpha|$; discontinuity in $\nabla\alpha$ (see text); ground truth region subscripts, 0 (black) to 5 (white). *Second row*: traslating cube; rightmost image represents part of $c(t)$ at intermediate, equispaced instants during the exposure, reconstructed as iso-alpha curves (theorem 2).

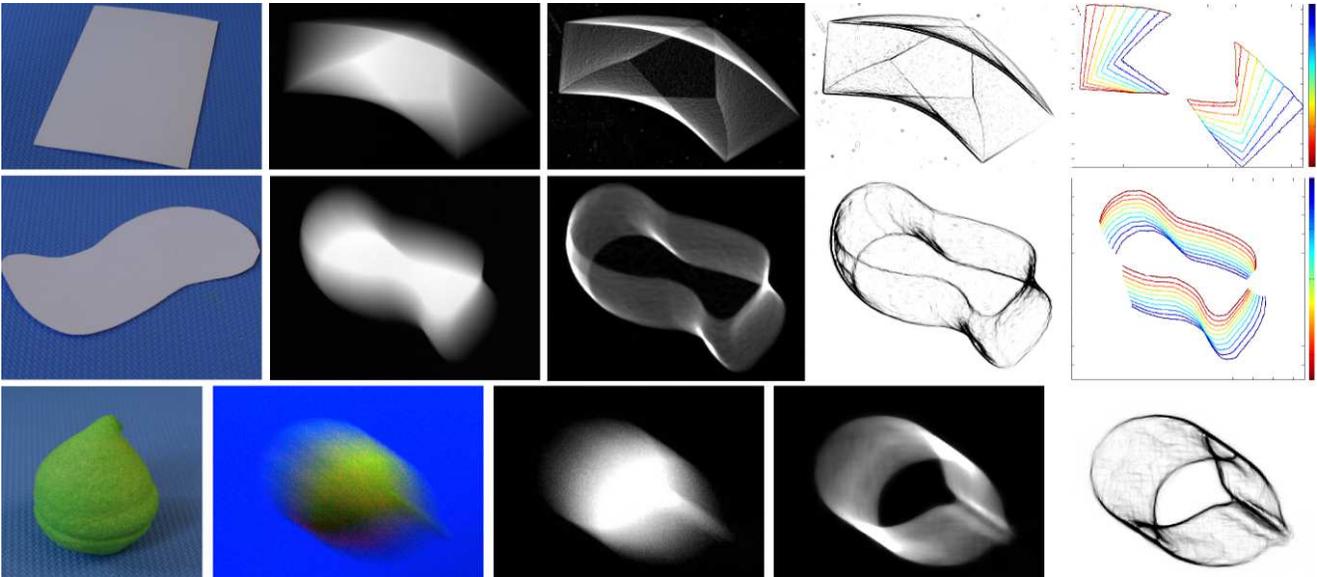


Figure 6. Results on real images. *First row*: white rectangular business card moving on a dark desk; from left to right: the object; the blurred image (from which α follows immediately); $|\nabla\alpha|$; discontinuity in $\nabla\alpha$; contour motion as iso-alpha curves. *Second row*: A planar object with curvy contours moving on a dark desk. *Third row*: A candy. Second image is the actual blurred image (exposure time 20 ms); note severe shading and noise. Alpha (third image) is computed using Levin's matting algorithm.

sophisticated application of the same principle can theoretically allow partial 3D reconstruction of a rototraslating object from a single image ([6]).

Theorem 5 allows to identify $c(t')$ and $c(t'')$, which is a valuable information in many practical applications; in figure 8, we exploit this in order to reconstruct the 3D position *and velocity* of a moving ball from a single slightly blurred perspective image.

Information about contour envelopes and corner paths can also be valuable for reconstructing motion in simple scenes, which has applications in machine vision or for understanding high-speed dynamic events with the high resolution and low price of an ordinary digital camera.

5. Experimental Results

The theoretical results we presented in the previous sections have been validated by experimental results both on real (Figures 6 and 8) and synthetic (Figure 5) images, in which ground truth is known.

Synthetic images are generated using the popular Blender raytracing software, by creating an animation with several hundred frames and averaging all together to obtain the actual motion-blurred image. Knowledge of the single frames also allows to automatically compute the ground truth for the region labeling, as in top right of figure 5. We extracted alpha data by means of the technique in [7], and verified that the result exactly matches the true alpha, gen-



Figure 4. Original image (left), and discontinuities computed directly over its gradient, without recovering α (right). As explained in Section 3.4, you can see the objects' contours at the beginning and end of the exposure, as well as the ball's contour at the bounce position (sudden speed discontinuity).

erated by rendering the object in each frame as plain white on a plain black background.

Real images have been shot using a digital reflex camera which allows manual control of exposure time; we both imaged uniform-color objects on an uniform background (which trivially translates to alpha data, possibly up to a scale factor if no $\alpha = 1$ region is visible), and objects meeting the color constraints in [7] over a known background. We used RAW images in order to avoid nonlinearities in the camera response function.

As a measure of $\nabla\alpha$'s discontinuity, we computed the norm of the derivatives of $\nabla\alpha$ along the x and y directions, then combined them to a single measure, which is plotted as the third column in figures 5 and 6; we implemented some basic median and gaussian filtering in order to mitigate the effect of noise and get recognizable shapes. We expect dramatic improvements by using more sophisticated techniques for finding discontinuities in $\nabla\alpha$.

6. Conclusions and future works

We presented a number of criteria to analyze motion-blurred object smears, and relate their features to the object apparent contour during the exposure period.

We initially isolate the contribution of the actual object motion by extracting the alpha matte of the semitransparent smear. Discontinuities in $\nabla\alpha$ partition the image into regions. We classify each region by means of two orthogonal criteria:

- the number of times each point in the region is crossed

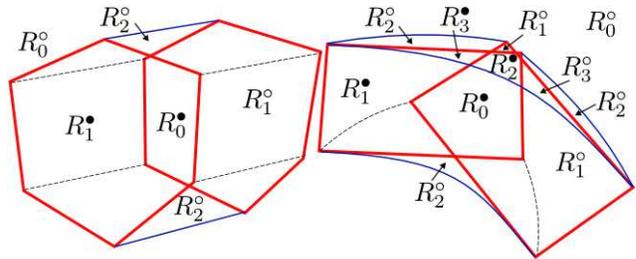


Figure 7. interpretation of cube (second row of Figure 5) and business card (first row of Figure 6) smears. Blue lines: type 2 region borders; thick red lines: type 1 region borders; thin dashed lines: paths of corners on the apparent contour (*not* a region border).

by the apparent contour during the exposure;

- whether the region is inside the object contour when the exposure begins.

Inside the regions crossed once by the apparent contour, iso-alpha curves represent the object contours at intermediate times during the exposure.

The object contours at the beginning and the end of the exposure are visible as region borders. The envelope of the object apparent contour over time also originates a region border. Discontinuities in $\nabla\alpha$ not mapping to a region border are still informative, since they either represent the path of a corner in the apparent contour, or the contour of the object when a sudden change in its motion occurs.

Experimental results confirm the validity of the theory we presented on both synthetic and real images. Our current work is aimed at:

- tuning existing matting algorithms for improved alpha matting of blurred smears, by exploiting the strong constraints on the smear color and alpha;
- improving and validating the experimental applicability of the current theory to temporal superresolution from an image or video; we plan to apply specific algorithms for region border detection, and implement automatic region labeling;
- studying how the combination of motion and defocus blur affects our theory.

References

- [1] N. E. Apostoloff and A. W. Fitzgibbon. Bayesian video matting using learnt image priors. In *Proc. of CVPR 2004*. 2
- [2] 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. 2
- [3] Y. Chuang, B. Curless, D. Salesin, and R. Szeliski. A bayesian approach to digital matting. In *Proc. of CVPR 2001*. 2

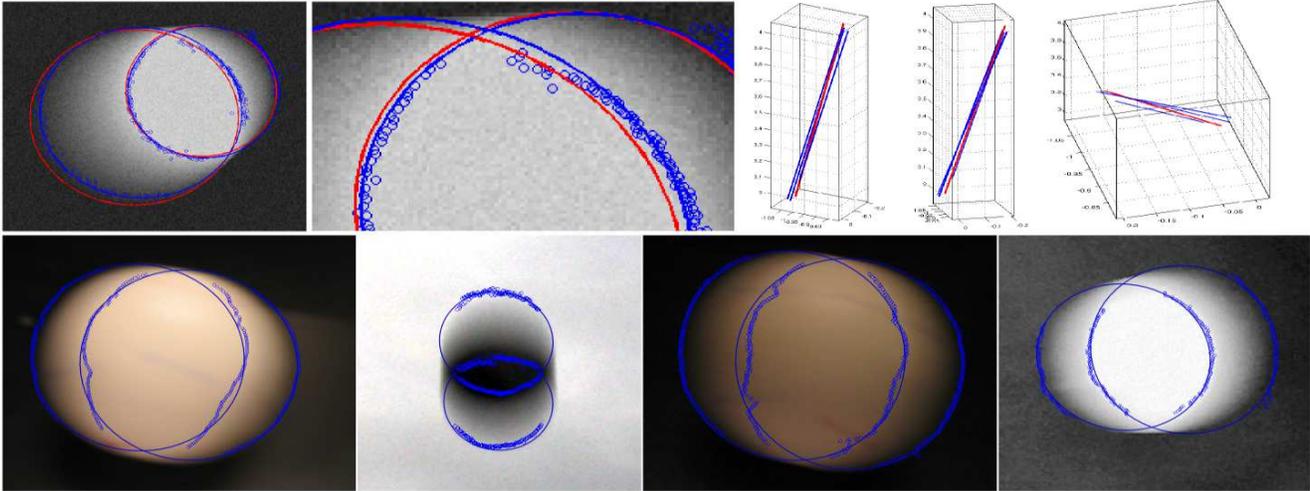


Figure 8. An application of our theory to 3D localization of a moving ball from a single perspective image. If the ball appears even slightly blurred, the traditional method of fitting an ellipse to its apparent contour does not work. We exploit theorem 6, and look for angular points along rectilinear intensity profiles parallel to the ball's motion; these points lie on $c(t')$ and $c(t'')$; by fitting two ellipses on the two sets of points, we are able to reconstruct the 3D position *and velocity* of the moving ball from a single image. Top row shows a synthetic example (with added gaussian noise): blue dots are discontinuities in the image intensity derivative along profiles, after denoising; red ellipses are the ground truth for $c(t')$ and $c(t'')$; blue ellipses are our estimate, obtained by fitting to blue dots; on the right, each 3D blue line represents our reconstructed position and velocity with a different realization of noise; red line is ground truth; other orientations of the same plot are also given. Bottom row shows the same algorithm on four different camera images (where ground truth is not available).

- [4] P. Favaro and S. Soatto. A variational approach to scene reconstruction and image segmentation from motion-blur cues. In *Proc. of CVPR 2004*. 1
- [5] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. In *ACM SIGGRAPH 2006 Papers*. 1
- [6] Y. Furukawa, A. Sethi, J. Ponce, and D. Kriegman. Robust structure and motion from outlines of smooth curved surfaces. *IEEE Trans. on PAMI*, 28(2):302–315, 2006. 6
- [7] A. Giusti and V. Caglioti. Isolating motion and color in a motion blurred image. In *Proc. of BMVC 2007*. 2, 6, 7
- [8] J. Jia. Single image motion deblurring using transparency. In *Proc. of CVPR 2007*. 1
- [9] S. Kang, Y. Choung, and J. Paik. Segmentation-based image restoration for multiple moving objects with different motions. In *Proc. of International Conference on Image Processing (ICIP) 1999*. 1
- [10] G. Klein and T. Drummond. A single-frame visual gyroscope. In *Proc. of BMVC 2005*. 1
- [11] A. Levin, D. Lischinski, and Y. Weiss. A closed form solution to natural image matting. In *Proc. of CVPR 2006*. 2
- [12] H.-Y. 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. 1
- [13] H.-Y. Lin and C.-H. Chang. Depth recovery from motion blurred images. In *Proc. of International Conference on Pattern Recognition (ICPR) 2006*, volume 1, pages 135–138. 1
- [14] H.-Y. Lin and C.-H. 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. 1
- [15] Y. Mishima. Soft edge chroma-key generation based upon hexoctahedral color space. U.S. Patent 5,355,174, 1993. 2
- [16] T. Porter and T. Duff. Compositing digital images. *Computer Graphics*, 1984. 2
- [17] M. Ruzon and C. Tomasi. Alpha estimation in natural images. In *Proc. of CVPR 2000*. 2
- [18] E. Shechtman and Y. Caspi. Space-time super-resolution. *IEEE Trans. on PAMI*, 27(4):531–545, 2005. Member-Michal Irani. 5
- [19] D. Slepian. Restoration of photographs blurred by image motion. *Bell System Tech.*, 46(10):2353–2362, December 1967. 1
- [20] A. R. Smith and J. F. Blinn. Blue screen matting. In *SIGGRAPH '96: Proc. of the 23rd annual conference on Computer graphics and interactive techniques*, pages 259–268. 2
- [21] J. Sun, J. Jia, C.-K. Tang, and H.-Y. Shum. Poisson matting. In *ACM SIGGRAPH 2004 Papers*, pages 315–321. 2
- [22] A. Yilmaz and M. Shah. Actions sketch: A novel action representation. In *Proc. of CVPR 2005*. 4