

Shadow Removal via Flash/Noflash Illumination

Cheng Lu, Mark S. Drew
School of Computing Science
Simon Fraser University
Vancouver, Canada V5A 1S6
{clu,mark}@cs.sfu.ca

Graham D. Finlayson
School of Computing Sciences
University of East Anglia
Norwich, U.K. NR4 7TJ
graham@cmp.uea.ac.uk

Abstract—Shadows bedevil multimedia applications, e.g. seeing into shadow regions in surveillance video. Model-based and non-model based, statistical methods, spatial, temporal, and invariant-based methods have been devised for combatting the shadow problem. Here we take a different approach, by attenuating the shadow by utilizing a second image under another illuminant to remove the effect of shadow-edges from an edge map of each frame. As a precursor step, we examine flash/noflash still image pairs. A flash image provides lessened shadows, but other shadows are produced. We can produce a flash-only (no ambient) image by subtracting the two images, but several artifacts remain. Instead, we have used the pure-flash image to detect the *ambient shadows* [ICME06]. However that method may fail when there are flash-shadows or specularities in the copy region. Here we manipulate the gradient field using a smoothing step including the directionality of edges near the shadow boundary, with improved results.

I. INTRODUCTION

Shadows bedevil multimedia applications, since shadows in video can cause object merging or object shape distortion. Aside from such tracking problems for tasks such as surveillance, seeing into shadow regions is important. Various model-based and non-model based, statistical methods, spatial, temporal, and invariant-based methods have been devised for combatting the shadow problem. Here we take a different approach, by eliminating the shadow (almost completely) by the mechanism of utilizing a second image under another illuminant to remove the effect of shadow-edges from an edge map of each frame. As a precursor step, we examine flash/noflash still pairs. The flash image provides lessened shadows, but other problems arise, and other shadows are produced. It is known that we can in fact produce a flash-only (no ambient) image by subtracting the two images, but the result is not as expressive as the ambient image, because of several artifacts due to the flash. Instead, we have previously used the pure-flash image to detect and remove the *ambient shadows*. Flash-edges are copied into the ambient-edge map, and re-integrated in the Fourier domain. However that method may fail when there are flash-shadows or specularities in the copy region. Here we introduce a smoothing step, including the directionality of edges near the shadow boundary. By manipulating the gradient field at the shadow edges, the difference caused by flash-shadow edges can be removed.

We argue that first going to a “spectrally sharpened” color space, and then focusing on the difference in a log domain of the flash image minus the ambient image, gives a very simple feature space consisting of two components — one in an illuminant-change 3-vector direction, and one along the gray axis. This space provides excellent separation of the shadow and nonshadow areas. Inserting edges from the flash image, within the ambient-shadow region, into the ambient image

edge map and inverting Poisson’s equation fills in the shadow. In addition, by manipulating the gradient field at the shadow edges, the difference caused by shadow edges can be removed. In this way, we arrive at an image with the advantages of the ambient-only image — warmth, no flash effects such as disturbing illumination dropoff with distance, pixel saturation etc. — but no shadows.

The combination of both ambient lighting plus a flash of course produces quite different illumination at a surface point than in an image taken under ambient lighting. For clarity, let us refer to the first image as “Ambient” and the second as “Both” (A and B). If we control the camera settings, or at least know them, and assuming there are no saturated pixels etc. [1], then $(B - A)$ should yield an image as if it were taken under the flash only (assuming one adjusts overall pixel magnitudes to compensate for camera settings, as in [2]). This is due to the fact that the B image consists of reflected light from the ambient sources plus from the flash. Lighting is from a different direction from the camera flash than the effective direction for ambient lighting (at each pixel). That is, there is a different visibility function for the flash image, and this produces “flash-shadows”. But since the pure-flash image, denoted “Flash” (F), sees reflected light from only the flash illumination, image F has no ambient shadows — no shadows that derive from the ambient, usually quite visible in both B as well as in A .

So we should be able to combine F with A , the ambient image, to be able to detect and hence eliminate the ambient-shadows. But we found that a simple differencing scheme does not work well: $F - A$ is large within the ambient-shadow region, but this is confounded by the change in color of the two illuminants: ambient and flash illuminations [3]. Here, we re-cast the difference image $F - A$ to make ambient-shadows easily detected. We go over to a color space in which a simple, diagonal, model of illumination change applies; then the difference of log images forms a feature space with two very simple components: a constant component along a 3-vector lighting-change direction, plus a difference of log intensity/shading/visibility terms along the gray $(1, 1, 1)$ axis. A projection into this plane leaves a large intensity difference between in-shadow pixels from A and from F , compared to out-of-shadow pixels. Results are seen to be very effective: ambient-shadows are simply identified.

Work on flash/noflash pairs started with [4] and then [2], [5], [6]. These works do not explicitly use the flash information for finding shadows (finding a “shadow matte”, in graphics terms—cf. [7]). Instead, these efforts are aimed at using the high-frequency information and detail from the lower-noise flash+ambient image to reduce noise in the ambient-only

image, e.g. using a joint bilateral filter. But the flash image by itself suffers from several drawbacks that make using the ambient-only image appealing: illumination dropoff with distance (producing a “tunnel effect” [8]), pixel saturation, and strong interreflections. The ambient image is known to be warmer and more pleasing.

A straightforward approach to change the ambient-lighting image inside shadows is to transfer the color information from a pure-flash image over to the ambient image. Methods to transfer color information between multiple images have been explored for various applications. Authors in [9] adopted a Bayesian framework to recover an image from a pair of images which have different exposure times. A color mapping function is generated by incorporating local spatial correspondence into global Gaussian color fitting. In [10], color transfer between two images is performed in a decorrelated color space $l\alpha\beta$, with better results than in RGB. However, these methods do not apply for our problem since we aim at not only retaining the ambient lighting effect but also finding and removing ambient shadows automatically.

Here we show that a simplified image model greatly aids in finding the shadows in the ambient image: naively using the intensity cannot succeed, since dark pixels can arise from many sources. But if we consider image formation, for the flash and ambient images, we can find a color space representation that separates the ambient-shadow feature vectors from the out-of-shadow locations. Our approach to removing ambient-shadows uses the high-quality gradient information from the flash image for the shadow regions and retains the nonshadow regions in the ambient image. To seamlessly combine these two parts, the image is reconstructed by solving Poisson equations.

In §II, a simple image-formation model is described that greatly simplifies the problem. Section III sets out how to promote the simple model via a color-space transform, and then examines the log-difference image for flash-no-flash pairs, showing that an illumination difference image can be generated with no surface-color term confounding lighting change. Section IV recovers the shadow mask for the ambient image, via projecting log-difference values onto the plane they mostly occupy and then detecting large differences by a robust procedure. Then §V shows how to copy flash information over to the ambient image, inside the shadow mask, and then re-integrate by combining gradients to guarantee integrability. The result is an ambient image with much less ambient shadowing.

II. IMAGE FORMATION

Let us restrict attention to Lambertian surfaces. Then at a surface point, under orthography, lighting is added up into a single effective light [11], taking into account visibility factors for each source. In general, the direction of the effective light is different in each color band [12].

Consider the RGB \mathbf{R} formed at a pixel \mathbf{x} , for illumination $E(\lambda)$ lighting a surface with reflectance function $S(\lambda)$. With 3 camera sensor sensitivity functions $\mathbf{Q}(\lambda)$, we have

$$R_k = \sigma \int E(\lambda) S(\lambda) Q_k(\lambda) d\lambda, \quad k = R, G, B, \quad (1)$$

where σ is Lambertian shading along with visibility.

We wish to go to a model that explains the difference in images formed under different lights by a simple diagonal 3×3 matrix. It has been found that this illuminant-change model holds, and in fact the image-formation description is greatly

simplified, if we make assumptions of Planckian lighting, Lambertian surfaces, and a narrowband camera [13]. In this case, we find that a log-difference image for flash F and ambient A images has a very simple form.

With a Dirac delta camera sensor $Q_k(\lambda) = q_k \delta(\lambda - \lambda_k)$, eq. (1) becomes simply

$$R_k = \sigma E(\lambda_k) S(\lambda_k) q_k. \quad (2)$$

Approximating lighting by Planck’s law, in Wien’s approximation:

$$E(\lambda, T) \simeq I k_1 \lambda^{-5} e^{-\frac{k_2}{T\lambda}}, \quad (3)$$

with constants k_1 and k_2 , temperature T characterizes the lighting color and I gives the overall light intensity. In this approximation, from (2) the RGB color R_k , $k = 1 \dots 3$, is simply given by

$$R_k = \sigma I k_1 \lambda_k^{-5} e^{-\frac{k_2}{T\lambda_k}} S(\lambda_k) q_k. \quad (4)$$

Define the following short-hand notations:

$$\begin{aligned} K &= \log(I k_1 \sigma); & s_k &= \log(S(\lambda_k)); \\ w_k &= \log(k_1 \lambda_k^{-5} q_k); & e_k &= -k_2 / \lambda_k \end{aligned} \quad (5)$$

Taking logarithms and showing retinal position \mathbf{x} , eq. (4) becomes

$$\log R_k(\mathbf{x}) = w_k + K(\mathbf{x}) + s_k(\mathbf{x}) + (1/T(\mathbf{x})) e_k \quad (6)$$

Here, the 3-vector w_k depends on the camera, as does e_k , independent of image location. The intensity and shading, encapsulated in K , do depend on location, as does the surface term s_k . Lighting temperature T depends on what lighting the surface point sees and adds up. *In shadows, both K and T are different* — the shadow region sees a different color and intensity of light than does the nonshadowed area. In this paper, we assume that lighting is effectively independent of location \mathbf{x} in each separate region: shadow and nonshadow.

Using these simplifying assumptions, we will use the fact that the ambient shadow is removed from the pure-flash image F to find the ambient-shadows in A .




Fig. 1. Sample images: (a): Ambient A , (b): Both B , (c): Flash F , (d): Flash shadows/specularities, (e): Ambient shadows, (f): Re-integrated image without using smoothing operation and (g): Re-integrated image via smoothing operation eq. (10).

III. IMAGE DIFFERENCE MODEL

A. Spectral Sharpening

The simplified model (4) is a better approximation if $\mathbf{Q}(\lambda)$ approximates a Dirac delta. We form an intermediate color space in which the sensors are optimally combined so as to

form new colors that better approximate color change induced by illuminant change via a diagonal model using Spectral Sharpening [14]. This applies a 3×3 transformation matrix M to the sensors, or directly to colors, so as to better enforce a diagonal model.

Since we mean to take logs, we need nonnegative colors from the camera data (with zero values treated specially). To do so, we carry out a ‘‘spectral sharpening with positivity’’ transform (cf. [15], [16]). Using calibration targets under two different lights, in [3] we find M via a novel optimization consisting of a constrained form of ‘‘database sharpening’’ (cf. [14]), but with hard positivity constraints: this forms M and a diagonal matrix D to minimize $\sum [M D M^{-1} A - B]^2$.

B. Log Image Difference

From (6), we notice that a difference image in log space can remove both a camera term w_k as well as the surface term $s_k(\mathbf{x})$. Let us form a *ratio image* by a difference image in log space: subtracting eq. (6) for two images, A and F , we have

$$\log R_k^A(\mathbf{x}) - \log R_k^F(\mathbf{x}) = [K^A(\mathbf{x}) - K^F(\mathbf{x})] + [1/T^A(\mathbf{x}) - 1/T^F(\mathbf{x})]e_k \quad (7)$$

for the difference between log pixel values under light A and light F , at pixel indexed by \mathbf{x} . Notice that the *surface term is entirely removed*, leaving a type of intrinsic illumination difference image. In this image, shadowing will stand out.

The ratio image arises from: (1) the intensity difference (with shading/visibility), proportional to the basic direction $\mathbf{u} = 1/\sqrt{3}(1, 1, 1)^T$; (2) a term proportional to the camera-dependent lighting-change 3-vector e_k . Notice that eq. (7) is pixel-wise: the two parts would be different inside a shadow, if one is present, since the lighting color temperature would be different there than in the rest of the image, as would the intensities. For details, including plots showing that eq. (7) actually obtains in real imagery, please see [3], [17].

IV. SHADOW-MATTE ALGORITHM

To find ambient-shadows, we need to find a plane in which most (non-shadow) pixels live, with basis vectors given by the $(1, 1, 1)$ direction \mathbf{u} and the lighting-change direction e_k , as in eq. (7).

It is easy to find this plane, using SVD. Then we can best model the log-difference by projecting onto this plane:

$$\mathbf{d} = \mathbf{V} (\log \mathbf{R}^A - \log \mathbf{R}^F) \quad (8)$$

where \mathbf{V} is the 2×3 pair of principal component vectors in the plane: \mathbf{d} is 2-dimensional. Projecting discards some information, but mostly discards noise.

We found that most of the variation in log-difference magnitude is along the \mathbf{u} direction (in the sharpened color space). Therefore we apprehend shadow formation as a process that generates large values in the log difference between Flash and Ambient (i.e., the \mathbf{u} direction), as well as differences in lighting color (the e direction). To take both into account we simply use the Euclidean magnitude in the plane. Since the logarithm is monotonic, and we are seeking large differences, this will find outlier values from both causes.

But first we must consider pixels that are in flash-shadows and specularities. We can regress between A and B to scale color, but cannot do so faithfully in saturated or specular areas:

these not only degrade the detail structure in highlight regions, but also generate abrupt changes in the flash image. In F , shadow pixels are entirely black (or at least one or more color channels is very dark). Fig. 1(c) shows an F image obtained from A (Fig. 1(a)) and B (Fig. 1(b)). Both A and B have specularities on the left side of the person’s face and on the wall, giving rise to saturation in the images. The saturated pixels only occur in out-of-shadow regions, but our shadow-removal task does need to locate these pixels.

We use a robust method [18] (Least Median of Squares–LMS) to find flash-shadows and specularities from $\|F\|$, using a 1D location mode-finder. The LMS gives outliers automatically, and these are shown in Fig. 1(d). In general, outliers will be either in flash-shadows, or in specularities, with outliers on the dark side of the mode and specularities on the bright side. We include both kinds of outliers in a ‘‘flash-shadow/specularity’’ mask since we wish to exclude them both.

Then we derive a magnitude \mathcal{I} in the plane as $\mathcal{I} = \|\mathbf{d}\|$. Again we use the LMS, and the resulting ambient-shadow mask (excluding the flash-shadows) is shown in Fig. 1(e). As an auxiliary step, we additionally produce a cleaner shadow mask by applying a Mean-Shift segmentation [19] to the \mathcal{I} image, before finding outliers.

V. SHADOW-FREE AMBIENT IMAGE RECOVERY

Finding the ambient-shadow mask takes us most of the way toward recovery of a *shadowless* ambient image. To proceed, we first map flash-image pixels to pixels in the ambient image, inside the shadow mask. To do so, we first regress on pixel values not in either the ambient shadow or flash shadow/specularities, from F to A . Then we use the resulting 3×3 matrix T to take flash pixels over to ambient ones. Then a straightforward approach would join the regressed F (in ambient shadow) and A (out of ambient shadow).

However, this does not entirely eliminate the ambient shadows. This is because (i) the regression does not give perfect color mapping from F to A ; and (ii) A has illumination discontinuities that do not usually exactly coincide with the mask boundary. But we may note that *edges* are not discontinuous in each image F and A inside the ambient-shadow mask. Therefore, to seamlessly join the two parts, we can blend the image edges from F and A by combining their gradient fields: if S is the shadow mask, then we define a new field $\log C$

$$\begin{aligned} \nabla \log C[S] &= \nabla (\log F[S]T) \\ \nabla \log C[\neg S] &= \nabla (\log A[\neg S]) \end{aligned} \quad (9)$$

VI. GRADIENT FIELD AT THE SHADOW BOUNDARY

But this gradient blending process gives a sharp switch in the gradient field at the shadow boundary. Now, we observe that shadow edges in fact are characterized by more or less gradual changes, and do not usually exactly coincide with the shadow mask boundary. Thus shadow effects may still remain in the later recovery process. We begin by simply dilating the shadow mask into nonshadow regions and then copy flash edges into the ambient edge map. The flash gradients contain no shadow edges, so this removes most shadow effects.

However this simple operation may fail in case there are flash shadows/specularities in the dilated region. As mentioned, the image F may be damaged in shadow/specularity regions and so flash gradients cannot be used in the dilated regions. Hence we encode a *smoothness step* for the ambient

gradients in the dilated regions, in order to simultaneously remove the difference caused by shadows and retain surface details. We take two factors into account: (i) the gradients of shadow edges are perpendicular to the nearby shadow boundary, and (ii) the characteristic of gradual shadow edges makes the Laplacian close to zero. Let Ω be the dilated region; we set the gradients $\nabla(\log \mathbf{C}[\Omega])$ in Ω via

$$\left\{ \begin{array}{l} \nabla(\log \mathbf{F}[\Omega] \mathbf{T}) : \text{ If } \mathbf{F}[\Omega] \text{ contains no shadows/spec.} \\ 0 : \text{ If } \Omega \text{ contains F shadows/specularities, and} \\ \\ \nabla(\log \mathbf{A}[\Omega]) \perp \text{ nearby shadow boundary, and} \\ \frac{1}{4} \nabla^2(\log \mathbf{A}[\Omega]) < \tau \\ \\ \nabla(\log \mathbf{A}[\Omega]) : \text{ Otherwise} \end{array} \right. \quad (10)$$

Taking another gradient and inverting Poisson's equation to recover $\log \mathbf{C}$ will blend the two gradient fields. We can simply use a Fourier transform on the entire image to re-integrate, a very fast method, using homogeneous Neumann boundary conditions. As well, in order to combine the Ambient and Flash gradients such that integrability is preserved, we project the combined set of edges onto an integrable convex set in the Fourier domain, before re-integrating via inverse Fourier transform [20], [21]. The re-integrated ambient image for our example image is shown in Fig. 1(f). Fig. 1(g) illustrates that our smoothing operation eq. (10) produces substantively better results than omitting it. Other images we have tried (space does not permit their display [17] but a thorough experimental evaluation was carried out) are also improved using the extra smoothing step eq. (10).

Altogether, the algorithm for ambient-shadow removal is as follows:

Camera calibration \rightarrow sharpening matrix \mathbf{M} .
 Transform to sharpened color space:
 $\mathbf{A}^\sharp = \mathbf{M} \mathbf{A}$; $\mathbf{B}^\sharp = \mathbf{M} \mathbf{B}$, $\mathbf{F}^\sharp = \mathbf{B}^\sharp - \mathbf{A}^\sharp$.
 Form logs, and log-difference $\log \mathbf{F}^\sharp - \log \mathbf{A}^\sharp$.
 Best 2D plane for log-difference feature \rightarrow basis \mathbf{V} .
 Project log-difference into 2D \mathbf{d} : eq. (8).
 Find flash-shadow/specularity mask $S_F =$
 outliers $\text{LMS}(\|\mathbf{F}^\sharp\|)$.
 Apply mean-shift filtering to $\|\mathbf{d}\|$.
 Find ambient-shadow mask $S_A =$ outliers $\text{LMS}(\|\mathbf{d}\|)$.
 Combine masks: $S = S_A \cap (\neg S_F)$.
 Copy color from \mathbf{A} into \mathbf{F} , to set scale for flash image
 edge map:
 $\mathbf{P} = \text{LS}(F[\neg(S_A \cup S_F)], A[\neg(S_A \cup S_F)])$,
 where LS is least-squares regression; then $\mathbf{F}' = \mathbf{P} \mathbf{F}$.
 Dilate S .
 Form combined log-image edge map $\log \mathbf{C}$ via eqs. (9,10).
 In each color channel, we take another
 derivative in the Fourier domain via a phase shift,
 and then project onto an integrable edge
 map whilst undoing the Laplacian.
 The inverted Poisson equation is unique only up to an
 additive constant, which we fix by regressing back to
 the ambient image $\log \mathbf{A}$, in $(\neg S)$.
 Exponentiate to recover non-log image.

VII. CONCLUSION

We present a method of removing shadows from ambient images via the log-difference of pure flash and ambient

images. This image removes the surface-reflectance, yielding intrinsic illumination and shading information. In an ambient shadow region, the difference between the two is large, from an intensity difference or a lighting-color difference or both, so finding the shadow area is straightforward. Copying flash-image information over to the ambient image, flash and ambient information is blended such that the resulting integrable edge field generates a shadowless, ambient-lighting, result. We present an edge-blending smoothing step that substantively improves the re-integrated image. The method removes attached shadows or cast shadows. No careful image segmentation is required in that no harm is done by copying over extra shadow edges. But since shadows contiguous to objects are indeed found, we will in future consider linking cast shadows with the objects which cast the shadow. Many interesting issues of shadow-cue perception become available once automatic shadow detection is available, such as interpreting spatial relation of objects from the changing displacement of shadows in video data.

REFERENCES

- [1] R. Ramanath, W. Snyder, Y. Yoo, and M. S. Drew, "Color image processing pipeline in digital still cameras," *IEEE Signal Processing*, vol. 22, no. 1, pp. 34–43, 2005.
- [2] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama, "Digital photography with flash and no-flash image pairs," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 664–672, 2004.
- [3] M. Drew, C. Lu, and G. Finlayson, "Removing shadows using flash/noflash image edges," in *IEEE Conf. on Multimedia and Expo: ICME*, 2006.
- [4] J. M. DiCarlo, F. Xiao, and B. A. Wandell, "Illuminating illumination," in *Color Imaging Conf.*, 2001, pp. 27–34.
- [5] E. Eisemann and F. Durand, "Flash photography enhancement via intrinsic relighting," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 673–678, 2004.
- [6] A. Agrawal, R. Raskar, S. K. Nayar, and Y. Li, "Removing photography artifacts using gradient projection and flash-exposure sampling," *ACM Trans. Graph.*, vol. 24, no. 3, pp. 828–835, 2005.
- [7] Y.-Y. Chuang, D. B. Goldman, B. Curless, D. H. Salesin, and R. Szeliski, "Shadow matting and compositing," *ACM Trans. Graph.*, vol. 22, pp. 494–500, 2003.
- [8] A. Agrawal, R. Raskar, S. K. Nayar, and Y. Li, "Removing photography artifacts using gradient projection and flash-exposure sampling," *ACM Trans. Graph.*, vol. 24, pp. 828–835, 2005.
- [9] J. Sun, C.-K. Tang, and H.-Y. Shum, "Bayesian correction of image intensity with spatial consideration," in *ECCV 2004*, 2004, pp. 342–354.
- [10] M. Gooch, B. Reinhard, E. Ashikhmin, and P. Shirley, "Color transfer between images," *IEEE Comp. Graphics & App.*, pp. 34–40, 2001.
- [11] W. Silver, "Determining shape and reflectance using multiple images," Ph.D. dissertation, MIT, 1980.
- [12] M. Drew, "Robust specular detection from a single multi-illuminant color image," *CVGIP: Image Understanding*, vol. 59, pp. 320–327, 1994.
- [13] G. Finlayson and S. Hordley, "Colour constancy at a pixel," *J. Opt. Soc. Am. A*, vol. 18, no. 2, pp. 253–264, Feb. 2001.
- [14] G. Finlayson, M. Drew, and B. Funt, "Spectral sharpening: sensor transformations for improved color constancy," *J. Opt. Soc. Am. A*, vol. 11, no. 5, pp. 1553–1563, May 1994.
- [15] M. Drew and G. Finlayson, "Spectral sharpening with positivity," *J. Opt. Soc. Am. A*, vol. 17, pp. 1361–1370, 2000.
- [16] K. Barnard, F. Ciurea, and B. Funt, "Sensor sharpening for computational color constancy," *J. Opt. Soc. Am. A*, vol. 18, pp. 2728–2743, 2001.
- [17] M. Drew, C. Lu, and G. Finlayson, "Ambient shadow detection and removal via flash/noflash pairs," Simon Fraser University, Tech. Rep., 2006, <http://fas.sfu.ca/pub/cs/TR/2006/CMPT2006-05.pdf>.
- [18] P. J. Rousseeuw and A. M. Leroy, *Robust Regression and Outlier Detection*. Wiley, 1987.
- [19] D. Comaniciu and P. Meer, "Mean shift: A robust approach toward feature space analysis," *PAMI*, vol. 24, pp. 603–619, 2002.
- [20] T. Simchony, R. Chellappa, and M. Shao, "Direct analytical methods for solving Poisson equations in computer vision problems," *IEEE Trans. Patt. Anal. and Mach. Intell.*, vol. 12, pp. 435–445, 1990.
- [21] G. Finlayson, S. Hordley, C. Lu, and M. Drew, "On the removal of shadows from images," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 28, pp. 59–68, 2006.