# Path-Based Shadow Removal

Clément Fredembach and Graham D. Finlayson
University of East Anglia
Norwich NR4 7TJ, UK
{cf,graham}@cmp.uea.ac.uk

**Abstract**

It has been shown that by thresholding the image gradient at the location of shadow edges and then reintegrating, shadow-free images can be obtained. Unfortunately, the current methods are computationally expensive and also create artifacts in the reintegrated image. Our proposed method uses non-intersecting random paths (also called Hamiltonian paths) to allow for fast 1D reintegration. Because the artifacts are due to missing gradient information, we further improve the results by inpainting the detected shadow edges as to prevent the occurrence of unwanted artifacts.

**Keywords**
Shadow Removal, Random Walk

## 1. Introduction

Every time an object lies in the way of an illumination source, a shadow is cast. Those shadows are not simply variation in brightness levels, they are typically a different colour than the rest of the scene. In many computer vision applications such as tracking, scene analysis and object recognition, shadows are a nuisance and hamper algorithm performance. Shadow removal may also be desirable for cosmetic reasons: a large tree may cast a large shadow on an otherwise beautiful landscape picture or the presence of a shadow may result in a dynamic range which cannot be displayed.

Assuming that the boundaries of the shadows in an image have been approximately found (e.g. we use the method set forth in [1]) there exist methods for synthesizing a shadow free image. The basic approach involves differentiating the image, setting derivatives at shadow boundaries to zero and reintegrating. In [1] the reintegration is formulated as a 2D Poisson problem and in [2] as a simple path based integration. The latter approach has the advantage of being very fast whereas the former results in images which look better. However, in both cases the recovered images have artifacts and suffer from low contrast. We have found this to be due to imprecision in shadow location estimation and artifacts due to the integration method.

The new path based approach is based on three insights. First, that we should not estimate the pixel values for the masked shadow boundary regions during reintegration. Second that path based reintegration works best when shadow and non shadow regions are reintegrated apart from one another. Third, that the values of the boundaries between shadow and non-shadow regions (pixels under the mask) can be found by inpainting.

## 2. Path-Based Methods

Shadow removal involves thresholding selected gradients within an image and then reintegrating. Let $I$ be an image and $\nabla I$ its gradient. We can threshold the derivatives using a function $T(\nabla I)$ such that

$$\begin{aligned} T(\nabla I) &= 0 \text{ if } |\nabla I| < \theta \\ &= \nabla I \text{ otherwise} \end{aligned}$$

To recover $I$ from $T(\nabla I)$, one has to approximate the integral by a mean square method. This amounts to solving a Poisson equation of the form

$$\nabla^2 I = \mathrm{div}(T(\nabla I)) \tag{1}$$

Where $\nabla^2$ is the Laplacian operator $\nabla^2 I = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2}$ and $\mathrm{div}(T(\nabla I)) = \frac{\partial (T(\nabla I))_x}{\partial x} + \frac{\partial (T(\nabla I))_y}{\partial y}$

Poisson equation can be solved using inverse Fourier transforms while taking care of boundary conditions issues.

The message of eq. 1 is that shadow removal is hard, moreover this problem is also ill-defined as there are 2 derivatives per pixel but we wish to recover 1 brightness level. In [2] we argued that a 1D (path-based) reintegration might be preferable for complexity and precision issues. If the path fulfils the condition that is

goes through every pixel once and once only, the problem is well-defined as at every step we are using either $dx$ *or* $dy$ depending on the current direction of the path.
Going through every pixel implies that the shadow boundary will be crossed plenty of times. Unfortunately, by thresholding the image gradient at such locations, we effectively introduce an error in the reintegrated image. Those artifacts being visually disturbing, we decided to minimize the number of passages through the shadow edge. To do so, we had to create a class of Hamiltonian paths (a path that goes through all image pixels once) on the graph spanned by the image (see Fig.1). Allowing a single entry/exit per shadow component proved to be determinant in obtaining visually pleasing shadow-free images.

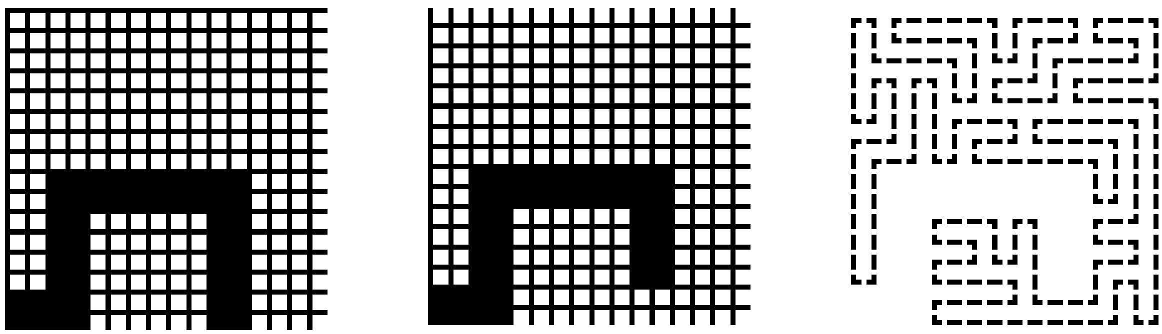

Figure 1: From left to right: an example of shadow mask (in black); the "opening" in the mask; a random Hamiltonian cycle used to visit every valid pixel.

### 3. Procedure

In order to obtain the shadow mask that will be used to distinguish between the shadow and non-shadow regions in the image, we proceed initially using the approach set forth in [1]. We first calculate a special grey scale invariant image where by construction shadows are not present (the image depends only on reflectance). Using this intrinsic image and the original colour image, we apply simple edge detection techniques to isolate the shadow areas of the original image. Basically we try to find which edges appear in the original but not in the intrinsic image (see illustration in figure 2).

Figure 2: The original image (left) and the corresponding intrinsic image (middle) and shadow mask (right).

While the approach works well using standard edge detectors such as SUSAN [4], the resulting shadow masks are not exact. Indeed, there are often many openings across what should be closed shadow boundaries. Openness is not suitable for the path-based integration because of the propagation effect (all information present at a pixel $i$ will be carried over until the end of the path and each time we go through an opening we effectively propagate incorrect information). The first contribution of this work is to close the shadow. We do this by using the initial SUSAN shadow mask estimate as a guide to finding shadow regions in a coarse region based segmentation provided by the mean shift algorithm [5]. For each edge given by SUSAN, we find the corresponding meanshift region. When an edge ends but is not closed, we select the boundaries of the meanshift region having the most included pixels and follow its boundaries until closure. For a large set of images this approach has proved to be consistently reliable.
Let $I_S$ the edge map obtained by SUSAN and $I_M$ the one obtained by mean-shift, let $e_S$ and $e_M$ further be the

edges in those maps. Let $N$ be the number of regions in the mean shift image. We then have

$$\forall e_S: \quad \alpha_{ij} = 1 \quad \text{if } e_S(j) \in e_M(i);\ i \in [1, N] \tag{2}$$

$$\alpha_i \quad = \quad \sum_j \alpha_{ij} \tag{3}$$

When an open point -no neighbor, yet not along the image boundaries- is encountered in $I_S$, we check which regions of $I_M$ are concerned (the ones having connecting edges). Among those regions, we select the one for which $\alpha_i$ is the greatest and continue the edge until closure. An illustration of the different edge maps and resulting closed edge can be found in figure 3.

Figure 3: SUSAN output (left), mean-shift segmentation (middle) and the resulting closed mask (right).

Once the closed mask is obtained, we allow one (random) opening per shadow component (there might be more than one) and start the reintegration from a non-shadow pixel (apart from the one path entry/exit to the shadow the two types of regions are integrated apart from one another). Applying this simple strategy, we find we can reintegrate shadow-free images.

The final step of the procedure is to deal with the shadow mask pixels themselves. As highlighted in figure 1, most shadow edge pixels will not be visited and will hold no information. To fill in the missing information, we use the technique described in [6]. The outline of this method is as follows. First, we compute all possible $11 \times 11$ windows for which all pixels are defined (not in the mask), let $N$ be that number. Since we are working in RGB, each window has a size$11 \times 11 \times 3$. Then, for each of the shadow mask pixels, we use a centered $11 \times 11$ window and compute its euclidian distance with respect to the $N$ $11 \times 11$ windows in the image. The window corresponding to the minimal Euclidian distance is then use to "fill in" the missing values. That is, we directly copy the pixels values of the chosen window at the "blank" pixels location. The procedure is then repeated until there are no more missing pixels. An example can be seen in figure 4, where the black portion of the image is the shadow mask.

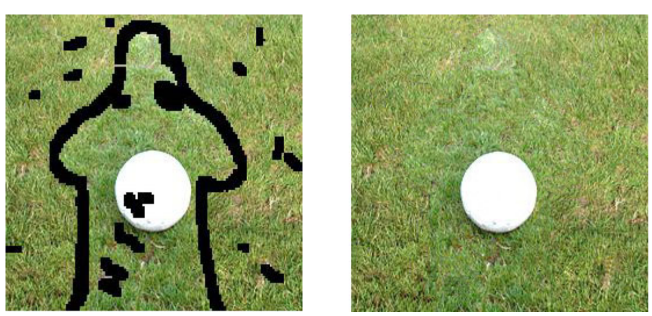

Figure 4: The recovered shadow-free image with the shadow boundaries (in black) on the left and the result of infilling on the right.

## 4. Results

The results we obtain using the described technique are generally robust and artifact as well as shadow free. Though a single path can work well, even better results are possible by averaging the results over a certain

number (say 4-5) of paths. This does not lead to a greater complexity as the most expensive steps (obtaining the mask and inpainting) are done only once per image regardless of the number of paths.
Figure 6 displays a variety of images obtained with our method. As one can see, they are generally successful with some exceptions. Some elements can perturb the reintegration, notably a non-exact shadow mask or the presence of colored noise in the shadow can alter the image gradient and impair the reintegration.

Figure 5: Results from our shadow removal method, all reintegrations are averaged over 4 paths.

**5. Conclusion**

To summarize, we have developed a method to reintegrate good quality shadow-free images from an original image and a shadow mask. We further proposed a framework to improve the resolution of such problems, notably the closure of shadow regions and the non integration of shadow edges.
To improve the robustness and speed of our method, we have to look at the weakest link, which for now is the shadow detection. This particular aspect appears to be the bottleneck of both quality and pace. We are currently investigating novel methods for fast and robust shadow detection that aim to solve this problem.

**Acknowledgments** G. Finlayson gratefully acknowledges the support of the Leverhulme Trust

**References**

[1] G. D. Finlayson, S. D. Hordley, and M. S. Drew. Removing Shadows from Images, *proceeding of ECCV 2002*, pp. 823–836, 2002.

[2] G. D. Finlayson and C. Fredembach. Fast Reintegration of Shadow Free Images, *Proc. of the IS& T 12th color Imaging Conference*, 2004.

[3] G. D. Finlayson and S. D. Hordley. Colour Constancy at a Pixel, *Journal of the Optical Society of America*, vol. 18, no 2, pp 253–264, 2001

[4] S. M. Smith. SUSAN - a new approach to low level image processing *Internal Technical Report TR95SMS1, Defence Research Agency* available at www.fmrib.ox.ac.uk/ steve, 1995.

[5] D. Comanicu and P. Meer. Mean shift: Arobust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Machine Intell*, vol 24, pp 603-619, May 2002.

[6] A. Criminisi, P. Perez and K. Toyama. Region Filling and Object Removal by Exemplar-Based Image Inpainting. *IEEE Trans. on Image Processing* vol 13, pp 1200–1212, Sept. 2004