

# CMPT 828: A shadow removal implementation

Ibraheem Alhashim  
Simon Fraser University  
School of Computing Science  
Burnaby, Canada  
iaa7@sfu.ca

## Abstract

*In this report we will discuss a simplified implementation of a shadow removal method. The key idea of the method is entropy minimization of the input image. An illuminant invariant image is computed and analyzed in order to extract shadow edges. The shadow edges are then feed into an image re-integration scheme that produces a colored intrinsic image that is free of dominant shadow. We discuss the different components and parameters and some of the challenges in practice. A number of outdoor results are presented that demonstrate the effectiveness of our simple implementation.*

## 1 Introduction

For different applications in computer vision, such as object tracking or recognition, it is often the case that failure is due to a visual discontinuity in the image. One major source of discontinuity are hard shadows that are typically prominent in outdoor scenes. Shadows in images or videos often interfere with the results of different algorithms utilizing edge or shape detection. Thus, a preprocessing or shadow filtering approach is needed for the effectiveness of some vision algorithms.

In this paper we report on a simplified implementation of the shadow removal method done in [3] and [5]. The method involves computing an illuminant invariant image, taken by an unknown camera, and identifying edges resulting from light changes in order to reconstruct a shadow-free image.

The following sections of this report are as follows: (a) we briefly describe some background and the assumptions considered in the method (b) we then describe the different components of the shadow removal algorithm (c) next we discuss the different experiments conducted and describe some of the parameters used (d) and lastly we show some results from our implementation.

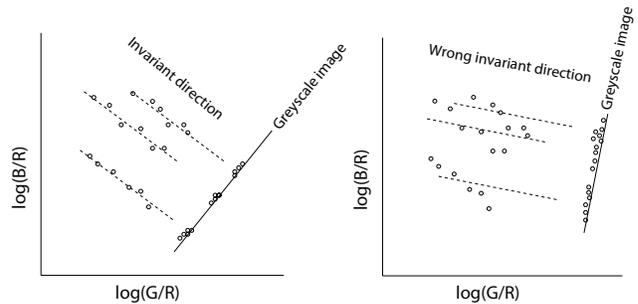


Figure 2: Two different projections during the process of entropy minimization. Left shows best direction for projection. Figure from [3].

## 2 Background

In order for this method to work, we first need to consider some basic assumptions about the input images being processed. Let us assume that the camera used when acquiring the images is a narrow-band camera, with or without spectral sharpening [6], with three sensors for each color channel, Red, Green, and Blue. Another assumption about the illuminant is that it belongs to the class of Planckian lighting. In all of our tests we consider sunlight and skylight to be the only sources of light. We also ignore specular highlights and other surface properties that do not appear in Lambertian surfaces. With these assumptions in mind, it is shown in [3] that the log-chromaticity (an intrinsic quality of a color) exhibit a linear behavior with lighting change. A consequence of this linear behavior is our ability to form an invariant image by projecting a 2D log chromaticity for each channel into a direction that best represent change in lighting.

For a calibrated camera with carefully constructed images it is possible to analytically compute the lighting invariant direction on which to project. For example, in [3], a color chart of 24 reflectance under 14 different lighting conditions are captured using a HP 912 digital camera with all

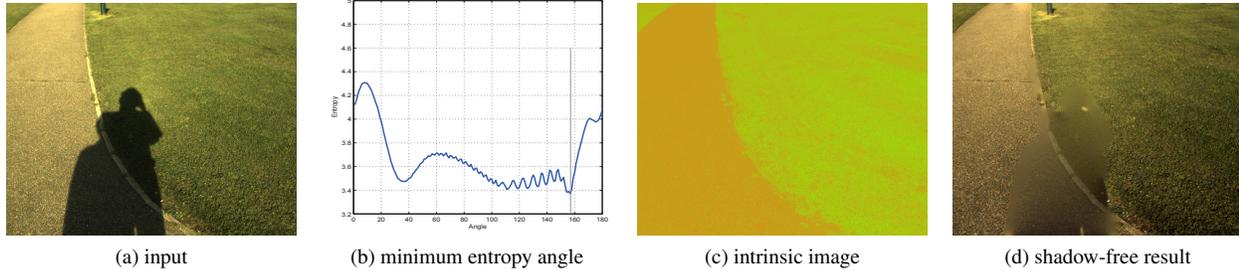


Figure 1: Shadow removal algorithm overview.

post processing disabled. The problem of finding the best projecting direction is then as simple as applying Singular Value Decomposition. However, in practice it is usually the case that such a calibration is not possible.

The suggested solution [3] to the missing calibration problem comes from the concept of entropy minimization. The key idea is that the best projection direction is the one with minimum entropy. In Figure 2, we can see two results of the pixel color projections in the 2D log-chromaticity space. If the greyscale image resulting from the left projection represent the lowest entropy it would form an intrinsic image with no shadows. Such images are then used to extract shadow lines that are then processed to produce our desired shadow-free images.

### 3 Shadow removal

The shadow removal algorithm we implemented consists of four main stages. The first stage is entropy minimization that takes the input image and returns a best projection angle with minimum entropy. The second stage is computing, using the found angle, the invariant  $L_1$  chromaticity image that is as shadow-free as possible. In the third stage, edges of both the source image and the shadow-free image are extracted and subtracted in such a way that only shadow edges remain. The final stage takes the input image and the extracted shadows and then modifies the image gradients such that when we integrate we get a shadow-free image. Figure 1 shows the three main stages, namely, entropy minimization, intrinsic image computation, and image reconstruction.

#### 3.1 Calibration & Entropy minimization

For images taken with calibrated known cameras it is possible to analytically compute the best projection direction that produces the best illuminant invariant image. However, we are interested in the problem of finding the best direction using only statics from the image itself. We apply the entropy minimization method described in more details

in [3]. We first compute a representation of the 2D log-chromaticity of the image. We then try to project for all angles  $\theta = 1 \dots 180^\circ$  and compute the entropy for each. Finally we select the best angle as the one resulting in the minimum entropy.

#### 3.2 Illuminant invariant images

The projected grey-scale invariant image resulting from the entropy minimization step represent an intrinsic, light-free image. One way to recover some color information to the image is to bring back light by offsetting the pixels in log-chromaticity space along the light invariant direction. The method we used is based on the work in [2] where the chromaticities of the brightest 1% pixels are recovered by an offset. The process is further refined using an optimization in order to increase color accuracy. The resulting intrinsic image represent a close enough representation of the input colored image without the effects contributed by the lights (in our examples both sun and sky).

#### 3.3 Shadow edges identification

At this stage of the algorithm we are now ready to extract the edges resulting from light change, i.e. shadow edges. The key idea is to extract both types of edges, light produced or intrinsic, from the input image and then filter out the shadow edges by comparing against extracted edges from our illuminant invariant image.

We start this process by applying the Mean-Shift smoothing filter [1] in order to avoid edges resulting from noise and high frequency image details. We then apply the Canny edge detector on both the smoothed image and the invariant image as shown in Figure 3(a). These edges are further expanded by applying a dilation operation in order to include a larger area of the shadow edges, especially since edges are almost never a single pixel wide. Finally, the shadow edges are computed as the edges that are present in the input image but not in the invariant image, see Figure 3(b).

### 3.4 Image reconstruction

The final stage of the algorithm is to compute the image's  $x$  and  $y$  gradients, then discount the effect of the shadow edges, and finally re-integrate by solving the well known Poisson equation [5]. In our implementation we smooth the shadow edges and then multiply them with the gradients in order to diffuse any drastic changes caused by this change in gradients. An essential part of this process is to enforce integrability, we do so by applying the method suggested in [5]. The entire process is applied to each color channel separately. An example of the output of this stage can be seen in Figure 1(d).

## 4 Analysis

Here we discuss some of our experiments and the choices we made in our implementation. We used Matlab for all operations in the algorithm and we also used code provided by Kai Barthel<sup>1</sup> for the mean-shift filter.

For the calibration stage we found that different forms of the input image result in different projection directions. For example, we tested the calibration process with smaller versions, smoothed, segmented, and combinations of these operations and found that simply calibrating on a quarter sized image results in the best calibration direction.

In the shadow edges identification stage, we experimented with different pre-processing methods, including four smoothing filters, to find the most robust and gen-

<sup>1</sup><http://rsbweb.nih.gov/ij/plugins/mean-shift.html>

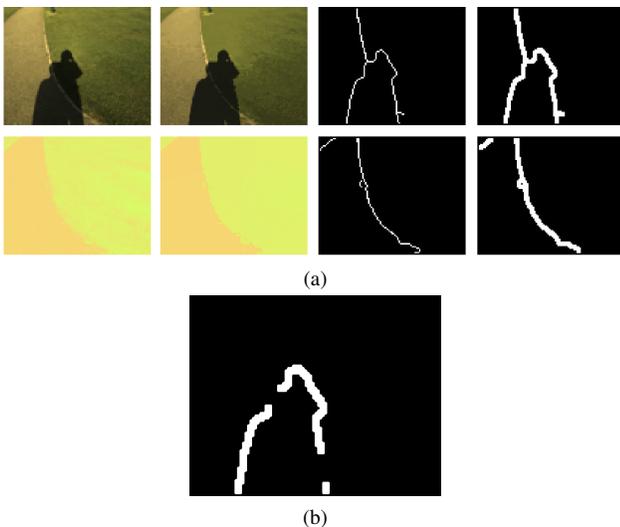


Figure 3: Shadows: (a) extracting shadows for (top) input image and (bottom) intrinsic image. (b) Only shadow edges remain.

eral approach. We found that mean-shift filtering did result in the least amount of spurious edges. Another parameter that is essential for good results is the Canny edge threshold which we eventually set in the range between 0.4 and 0.7. Also the dilation process we applied requires a parameter that specifies a width, we found that specifying it to match the width of the shadow border region results in the best reconstruction. However, this may also lead to larger smoothed areas in the output.

For the image reconstruction step we experimented with varying the thickness and effect of the identified shadow edges. We found that some level of blurring can help hide some of the visible artifacts near the shadow edges.

We tested JPEG compressed images and we did not obtain good results. This is mainly due to the error when computing the minimum entropy and the presence of outliers. Possibly, either noise removal procedures or selective data filtering might alleviate some of these errors.

## 5 Results

We present different results of shadow removal using our implementation. The average processing time for a 500px image is around 15 seconds. In some images the default parameters are modified to better suite the image's properties such as texture or complexity. A gallery of results is shown in Figures 4 and 5. The last example was taken with a modern Sony NEX-5N digital camera, however, it seems that it failed possibly due to the complexity of the scene.

## 6 Conclusions

We have implemented a basic version of the shadow removal method described in [3]. Some of our results show successful shadow removal even with our simplistic shadow detection approach. The main observation of this method to shadow removal is the reliance on entropy minimization to compute an illuminant invariant image that is free of most distinct shadows. Once such an intrinsic image is computed, we are able to discard the influence of shadows in outdoor scenes which can help in object tracking or recognition applications. Later work on the same topic uses a different formulation of entropy [4] that seems to suggested that the approach still holds promise for future improvements. Finally, we note that our current implementation is not able to handle many other cases due to the naive edge selection process and the simple re-integration procedure.

## References

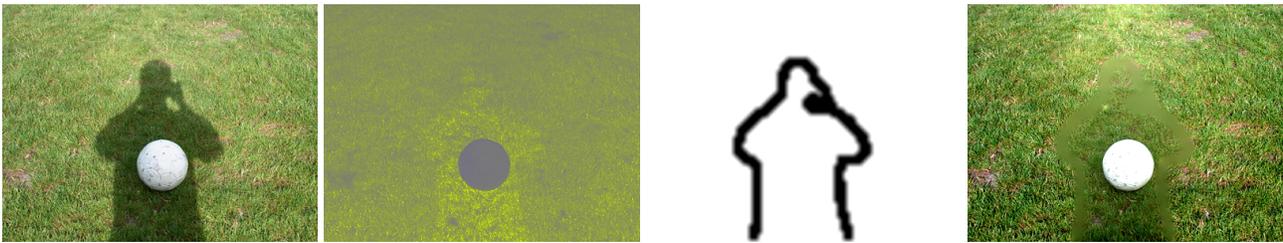
- [1] D. Comaniciu and P. Meer. Mean shift analysis and applications. In *Computer Vision, 1999. The Proceedings of the*

*Seventh IEEE International Conference on*, volume 2, pages 1197–1203 vol.2, 1999.

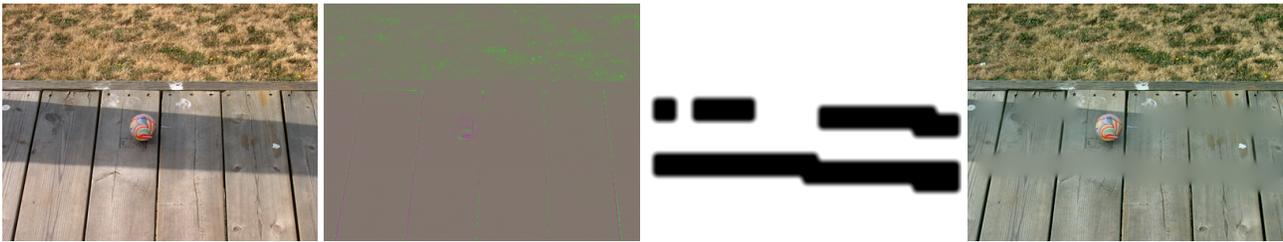
- [2] M. S. Drew. Recovery of chromaticity image free from shadows via illumination invariance. In *In IEEE Workshop on Color and Photometric Methods in Computer Vision, ICCV 03*, pages 32–39, 2003.
- [3] G. Finlayson, M. Drew, and C. Lu. Intrinsic images by entropy minimization. In *Computer Vision - ECCV 2004*, volume 3023 of *Lecture Notes in Computer Science*, pages 582–595. Springer Berlin / Heidelberg, 2004.
- [4] G. Finlayson, M. Drew, and C. Lu. Entropy minimization for shadow removal. *International Journal of Computer Vision*, 85:35–57, 2009. 10.1007/s11263-009-0243-z.
- [5] G. Finlayson, S. Hordley, C. Lu, and M. Drew. On the removal of shadows from images. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 28(1):59–68, jan. 2006.
- [6] G. D. Finlayson, M. S. Drew, and B. V. Funt. Spectral sharpening: sensor transformations for improved color constancy. *J. Opt. Soc. Am. A*, 11(5):1553–1563, May 1994.



(a)



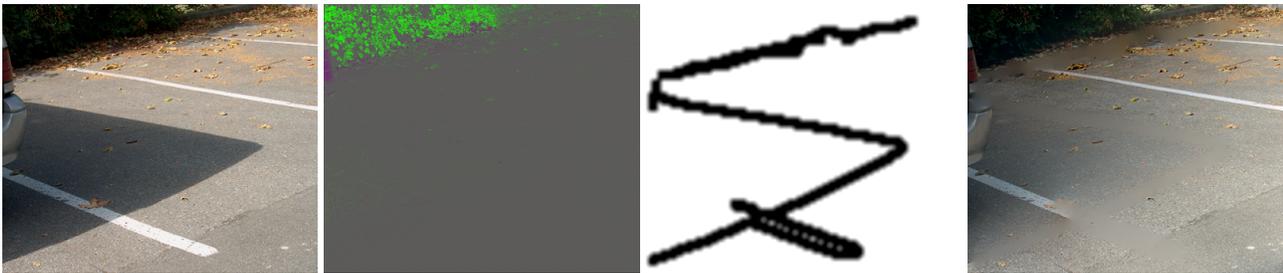
(b)



(c)



(d)

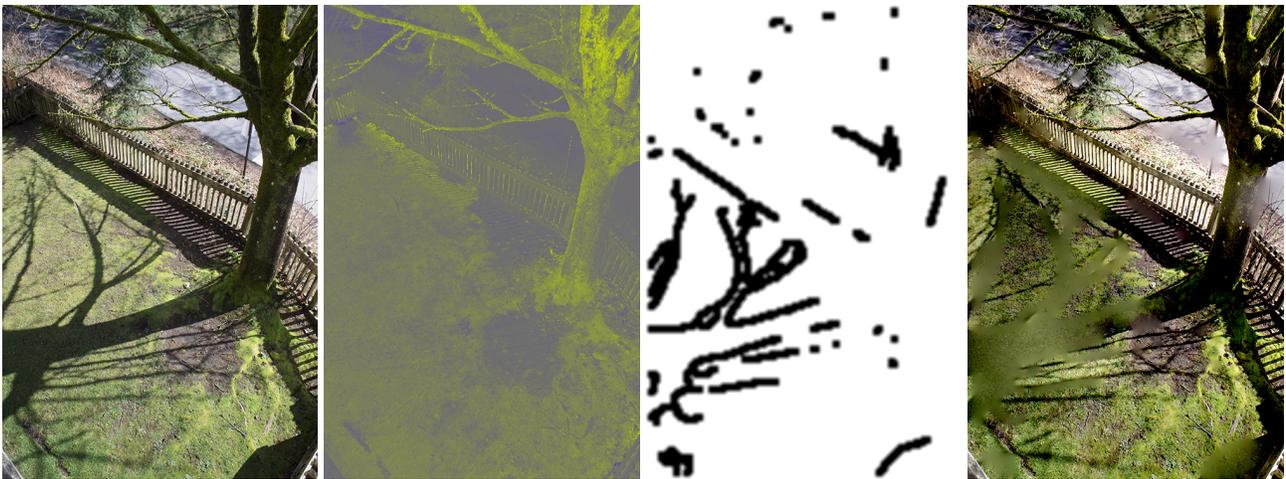


(e)

Figure 4: Results gallery.



(a)



(b)

Figure 5: Results gallery.