

OPTICAL FLOW FOR NON LAMBERTIAN SURFACES BY CANCELLING ILLUMINANT CHROMATICITY

Chetan Arora* Michael Werman

The Hebrew University of Jerusalem
Jerusalem, Israel

ABSTRACT

Optical flow, the pixel level correspondences between a pair of images is an important problem in computer vision. Standard optical flow computation algorithms assume constant brightness and fail on specular surfaces. Earlier work to alleviate problems with specular surfaces evaluate the illuminant chromaticity using a few correspondences in the images and then jointly optimize flow and appearance under the dichromatic model. We argue that the correspondences obtained by these methods are mostly pairs of pixels that are Lambertian thus giving a noisy estimate of the illuminant chromaticity. We suggest a new approach to evaluate the illuminant chromaticity which does not require exact correspondences and gives a better estimate of illuminant chromaticity. We use the evaluated chromaticity to project the input images on to a specular invariant color space and show that standard optical flow algorithms on this color space significantly improves the flow results. The suggested approach is simple, efficient and more importantly can utilize existing algorithms to compute optical flow on non Lambertian surfaces.

Index Terms— Optical Flow, Specular Surfaces, Non Lambertian Surfaces

1. INTRODUCTION

Finding optical flow is a crucial step for many computer vision problems. Many of these algorithms compare raw pixel intensity or color and assume these are the same at corresponding pixels. However, the constancy assumption is valid only for Lambertian surfaces. Under the dichromatic model for dielectric materials proposed by Shafer [1], light reflected from an object is a linear combination of diffuse and specular components. The diffuse component follows the Lambertian model and has radiance invariant to the viewing angle. The specular component expresses the directional reflection of the incident light hitting the object. While the color of diffuse component is the intrinsic property of the object's surface, the specular component is of the color of incident light, illuminant chromaticity. For non Lambertian surfaces the presence



Fig. 1: We consider the problem of estimating optical flow at highly specular surfaces. Regular optical flow estimation algorithms such as LK [2] assume constant brightness at corresponding pixels and fail on such surfaces.

of a viewing angle dependent specular component can result in large differences between the colors/intensities of corresponding pixels (see Fig. 1). As a result standard optical flow algorithms such as LK [2] relying on a error measure based upon constant brightness assumption may give invalid flow/correspondence. The flow algorithms working on intensity gradient such as SIFT are also likely to suffer since intensity gradients for a non Lambertian surfaces viewed from different angles are not preserved.

Optical flow algorithms for non Lambertian scenes assume pre-processing which can cancel the specular component of the observed pixel color/intensity by estimating and subtracting the illuminant chromaticity. Retinex theory [3] provides a computational scheme for color constancy. However the presence of various light sources, shadows and curved surfaces makes its application difficult in practice even for natural scenes. [4, 5] proposed methods to estimate illuminant chromaticity with various assumptions on the scene. These assumptions are problematic when dealing with highly textured surfaces and are not consistent with our requirements.

Yoon and Kweon [6] suggested a method to find correspondence under a white illuminant by computing a specular free two band image. Yang et al. [7] suggest to search for correspondence by analyzing the chromaticity of the color difference between two corresponding pixels in two images. The suggestions implicitly assume correct correspondence

*Chetan Arora is now with IIIT Delhi

at some pixels to arrive at a valid chromaticity hypothesis. We argue that most correspondences obtained by these methods are pixels/regions which are ‘almost’ Lambertian and therefore the constant brightness assumption ‘approximately’ holds. The specular component at these pixels is small and estimating the illuminant chromaticity from such pixels is noisy. Ideally a good estimate of illuminant chromaticity can be obtained from the region which are highly specular. But these are precisely the places where optical flow algorithms fail to find correspondence.

We observe that in many optical flow problems the two input views are coarsely aligned or it is possible to coarsely align them using correspondences at less specular points. The failure of an optical flow algorithm in a region is in itself an indication of specularity (or a no texture area which can also be easily identified). Since the views are coarsely aligned (due to the Lambertian pixels) the histogram of these regions is expected to be close according to the Lambertian assumption. Any divergence from an equal histogram observation is therefore due to the specular component which in our case expresses illuminant chromaticity. In this paper we suggest an iterative algorithm which uses chromaticity estimates from regions where the previous optical flow iteration fails. This chromaticity estimate is then used to find a new color space for the image with one of the color axis along the illuminant chromaticity. The input images are projected to this new color space for the next iteration of optical flow. We show that this improves by upto 50%, the number of pixels with correct optical flow in the experiments we conducted.

2. ILLUMINANT CHROMATICITY ESTIMATION

Under the dichromatic reflection model proposed by Shafer [1], the reflectance of a surface can be split into two components, surface body reflectance (diffuse) and interface reflectance (specular). The diffuse component follows Lambert’s law and has the same radiance when viewed from any angle. The specular component is due to the illuminant chromaticity and captures incident light reflected at a surface. The light from a surface is a linear combination of two components and can be expressed as:

$$I(p) = D(p) + \alpha(p)L,$$

where $I(p) = (I_r(p), I_g(p), I_b(p))$ is the observed color at pixel p , $D(p) = (D_r(p), D_g(p), D_b(p))$ is the diffuse and $L = (L_r, L_g, L_b)$ is the global illuminant color. Note that we assume a global illuminant L irrespective of pixel location. α is a scalar capturing surface properties, spatial position with respect to illuminant and local geometry of the scene surface.

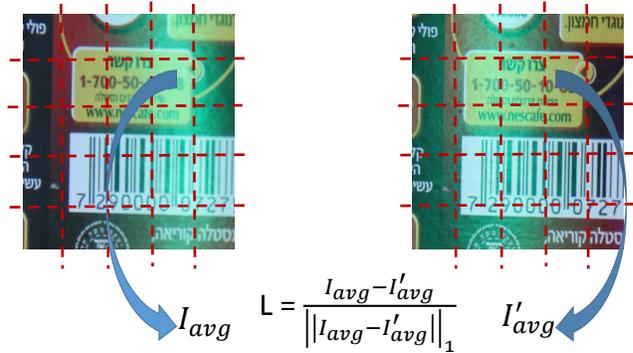


Fig. 2: We estimate illuminant chromaticity by taking the normalized difference of the averages of two corresponding regions. Regions are at the sweet spot between two extreme positions represented by corresponding pixels and the difference of a global histogram. The suggested approach allows us to choose specular regions yet use the similar histogram assumption without requiring exact correspondence.

2.1. Chromaticity Estimation from Corresponding Points

Observed colors $I(p)$ and $I(p')$ at a pair of corresponding pixels can be expressed as:

$$\begin{aligned} I(p) &= D(p) + \alpha(p)L \\ \text{and } I(p') &= D(p') + \alpha(p')L. \end{aligned}$$

Given that the corresponding pixels are images of the same 3d point, the diffuse component (D) is equal. The global illuminant can therefore be expressed as:

$$\begin{aligned} L &= \frac{I(p') - I(p)}{\alpha(p') - \alpha(p)} \\ L &= \beta(p) \frac{\Delta I(p)}{\|\Delta I(p)\|_1}, \end{aligned} \quad (1)$$

where $\|I\|_1 = I_r + I_g + I_b$ and β is a scalar. Every pair of corresponding pixels gives an estimate of illuminant chromaticity upto scale. Corresponding points at an ‘almost’ Lambertian surface leads to a small specular component and therefore a smaller difference between their observed colors. This leads to a noisy estimate for the illuminant chromaticity due to the denominator of Eq 1.

2.2. Chromaticity Estimation at Specular Regions

Assuming the images to be coarsely aligned, we expect their histogram and therefore average color to match under the Lambertian assumption. The difference in global average is therefore also an estimate of illuminant chromaticity. The global average is easy to compute without requiring pixel level correspondence. However, we expect any real scene to contain both surfaces which are highly diffuse as well as highly specular. A global average will therefore over-smooth the chromaticity estimate. On the other extreme of this hypothesis is the corresponding pixel scenario where the

estimate will be precise but the histogram will be the same only with exact correspondence, something which is hard to get, especially at specular regions.

We suggest a middle approach, not averaging over the whole image, but using regions that we can assume have similar histograms even without exact correspondence (see Fig. 2). The important question of which regions to choose for illuminant chromaticity estimation is also straightforward given that these will be the regions where standard optical flow algorithms fail due to the mis-assumption of constant brightness. Note that standard optical flow algorithms may also fail in textureless regions. We filter these regions from the selected ones based on the edge score in the region.

Once each region gives its estimate of illuminant chromaticity, we average (per color channel) all estimates to arrive at a global illuminant chromaticity. More robust statistical descriptors based upon median or clustering did not improve the performance significantly and therefore were not used in the experiments.

3. SPECULAR INVARIANT COLOR SPACE

One objective of this paper is to utilize existing optical flow algorithms even for specular surfaces. Therefore we would like to project the image to a color space where the constant brightness assumption holds and then use existing flow algorithms. Shafer [1] showed that reflection at a specular region is the sum of a view invariant diffuse component and illuminant chromaticity. Therefore we subtract the illuminant chromaticity component as follows to recover the viewpoint invariant diffuse component:

$$I_{\text{new}}(p) = I_{\text{orig}}(p) - (I_{\text{orig}}(p) \bullet L), \quad (2)$$

where ‘ \bullet ’ represents the dot product. We propose to use the new color values computed as above in the optical flow algorithm. Note that unlike original RGB space, components of the new color vector may be negative. The standard implementation of optical flow algorithms such as LK available in OpenCV [8] expect non negative RGB values (however this is not required for the theoretical working of the flow algorithm). We thus used our own implementation which can handle negative valued chromaticity components in our experiments.

4. PROPOSED ALGORITHM

We run Algorithm 1 iteratively, re-estimating the illuminant chromaticity at each step, and finding optical flow in the projected color space. In each iteration we progressively compute an improved optical flow, since the projected color space cancels out more accurate illuminant color. The better alignment allows us to do an improved average color matching in specular regions and thus improve the accuracy of illuminant

Algorithm 1 SpecularFlow Algorithm

- 1: **while** Optical flow improves **do**
 - 2: If not first iteration project image to color space perpendicular to estimated illuminant.
 - 3: Find optical flow using a standard optical flow algorithm (e.g. LK [2]).
 - 4: For regions where optical flow fails, assume zero (or neighbor average) flow.
 - 5: Find regions with significant difference in average color using the current optical flow estimation.
 - 6: Reject smooth regions.
 - 7: Evaluate estimated illuminant chromaticity using the difference of average colors in corresponding regions.
 - 8: Average the estimates obtained from the different regions to get a single global illuminant chromaticity.
 - 9: **end while**
-

color estimation. The process is repeated until there is no change in optical flow.

5. EXPERIMENTS

We conducted experiments on the images available from the Middlebury 2006 stereo dataset [9]. There were images available with multiple illuminations and exposures and we used view0 and view1 with random illumination and exposure. 870 points (shown in red) were sampled uniformly on the image and optical flow found on each of these points using LK [2] based on color difference. As described above, we used our implementation of LK which can work with negatively valued chromaticity components. We find the difference of average color in each of the regions (64×64) around the sampled and tracked point. We do not average saturated (>250) or dark (<5) pixels and reject regions where the L1 norm of the difference is less than 20. The rest of the regions are used for finding illuminant chromaticity. Table 1 gives the results. There is an improvement in flow in most of the examples we tested even though this database did not have much specular-ity in most images.

We also conducted experiments on our sequence where we controlled the illuminant chromaticity by illuminating a specular object with a green lamp. Fig. 3 shows the result. 900 points were chosen uniformly for finding optical flow. In the first iteration LK finds flow on only 102 of the 900 points. Many of the regions where optical flow was successful were not chosen for finding illuminant. With the remaining regions, the algorithm correctly identifies the illuminant as a mix of green and blue. While the lamp was green, blue can be attributed to ambient light. There is a 4 times improvement in optical flow after cancelling the illuminant chromaticity.

Fig. 4 shows the difference in projected input images (after aligning and warping one to the other) after iterations 1 and 2. Higher green values in the first image shows the extent of specular-ity in the images.

Input			Flow Success			Num Iters	Estimated Illuminant Chromaticity
Dataset	Illum	Exp	Without Correction	Proposed Method	Using Valid Tracked		
Cloth2	3	2	837	845	837	5	(0.364430, 0.304564, 0.331006)
Cloth4	3	2	783	812	783	6	(0.247378, 0.296470, 0.456152)
Baby2	3	2	771	800	771	7	(0.260307, 0.346704, 0.392989)
Rocks1	2	2	856	857	856	3	(0.371439, 0.337341, 0.291220)

Table 1: Experiments on images from Middlebury 2006 stereo dataset [9]. Images were randomly chosen from available illuminations and exposures. Num Iters and Estimated Illuminant chromaticity is from proposed method. Note that, we are not aware of the ground truth illuminant in these cases. The algorithm chooses a chromaticity which makes the corresponding regions most similar. As is our thesis, finding illuminant chromaticity using tracked regions fails to find any significant change and does not help improving flow in any of the tested samples.



Fig. 3: Optical flow of coffee jar illuminated with green light. (a),(b) Input images. Shiny label exhibits significant specularities. (c) Optical flow computed at regularly spaced grid points using LK [2]. Specularities create problems and flow computation succeeds only on 109 grid points out of 900. Blue dots indicate regions where the flow computation is successful. (d) Estimated illuminant chromaticity from each region with significant change in average color. We reject regions with saturated colors, dark regions and non textured regions which can give noisy estimates of the illuminant chromaticity. The global illuminant estimated is (0.090247, 0.539450, 0.370303). This is plausible since we illuminated the scene with a green light bulb. The high blue component can be attributed to ambient light. (e) Optical flow computed after projecting the image to a new color space. Flow computation is now successful at 361 points.



Fig. 4: (a) Difference in two input images (after ground truth alignment) for experiment corresponding to Fig. 3. The large difference in green channel shows the extent of specularities. Small differences at edges is inaccuracy in ground truth warping and should be ignored. (b) Difference after projecting to color space invariant to illuminant chromaticity.

6. CONCLUSION

We presented a method to improve optical flow computation for non Lambertian surfaces. Most existing methods rely on estimating illuminant chromaticity from regions where the optical flow algorithms succeed and use the change in observed color to estimate the illuminant. We showed that such regions often do not have significant changes in observed

color and are therefore not much help in estimating illuminant chromaticity. Instead using the difference of average of colors in sufficiently large windows gives a better estimate of chromaticity which can be progressively improved as we correct the illumination estimation and find optical flow in more regions. Experiments conducted on our as well as publicly available images validates the proposal.

7. ACKNOWLEDGEMENTS

This work was supported in part by the Israel Science Foundation and the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI).

8. REFERENCES

- [1] S. A. Shafer, "Using Color to Separate Reflection Components," *Color Research and Applications*, vol. 10, no. 4, pp. 210–218, 1985.
- [2] Bruce D. Lucas and Takeo Kanade, "An iterative image registration technique with an application to stereo vision," in *Proceedings of the 7th International Joint*

Conference on Artificial Intelligence - Volume 2, 1981, IJCAI'81, pp. 674–679.

- [3] Edwin H. Land, “The Retinex Theory of Color Vision,” *Scientific American*, vol. 237, no. 6, pp. 108–128, Dec. 1977.
- [4] Hsien-Che Lee, “Method for Computing the Scene-Illuminant Chromaticity from Specular Highlights,” *Journal of the Optical Society of America A*, vol. 3, no. 10, pp. 1694–1699, Oct. 1986.
- [5] Veronique Prinet, Dani Lischinski, and Michael Werman, “Illuminant chromaticity from image sequences,” in *The IEEE International Conference on Computer Vision (ICCV)*, December 2013.
- [6] Kuk-Jin Yoon and In-So Kweon, “Correspondence search in the presence of specular highlights using specular-free two-band images,” in *Computer Vision - ACCV 2006, 7th Asian Conference on Computer Vision, Hyderabad, India, January 13-16, 2006, Proceedings, Part II*, 2006, pp. 761–770.
- [7] Qingxiong Yang, Shengnan Wang, N. Ahuja, and Ruigang Yang, “A uniform framework for estimating illumination chromaticity, correspondence, and specular reflection,” *Image Processing, IEEE Transactions on*, vol. 20, no. 1, pp. 53–63, 2011.
- [8] G. Bradski, “Opencv ver 2.4.3,” 2013.
- [9] “Middlebury 2006 Stereo Dataset,” <http://vision.middlebury.edu/stereo/data/scenes2006/>.