

DETECTION OF SPECULARITIES IN COLOUR IMAGES USING LOCAL OPERATORS

Paul N.E. Pellicano
School of Computing Science
Simon Fraser University, Burnaby B.C.

ABSTRACT

Areas of spectral reflectance, or highlights, can be analyzed for a wide range of information or clues that they give us about a scene. This paper presents a local algorithm for analyzing a moderately unconstrained color image to determine the areas of spectral reflectance. The algorithm is based on the separability of the diffuse and spectral reflection components by differential methods.

The location of specular reflectances are marked by finding zero-crossings in concave down regions for two-dimensional arrays of intensities representing the color image. These zero-crossings correspond to the centers of the highlight regions. The highlight centers are then expanded to highlight regions by region growing in a direction orthogonal to the local orientation of the highlight. Thus, at the conclusion of the algorithm, the information known about each highlight includes location, size and direction.

INTRODUCTION

Computer Vision is often perceived as something that should be trivial. The reason for this perception is that we are ourselves so good at vision, we take the whole process of vision for granted. In fact, the interpretation of our three dimensional world, as portrayed in a two dimensional array of intensities, is anything but trivial. While humans bring a vast amount of 'intrinsic' information to bear on the problem of image analysis, the computer does not have the capacity, at the current time, to perform the same feat. Therefore, to permit any useful analysis of an image whatsoever, we tend to limit, or constrain our image world such that analysis becomes feasible with respect to the limited amount of knowledge we can impart to the computer.

A particularly useful and efficient task that we practice every day, however unwittingly, is that of distinguishing between objects made from different materials. An important prerequisite for such perception is the ability to discern the quality of an object's appearance. Various qualities of appearance are apparent in the world around us, such as texture, color, shine, luster, etc. all of which

give us important clues as to an object's composition. This paper will concern itself with the quality of surface gloss.

Glossiness, in general, is correlated with specular reflectance [Beck 72]. By looking at picture 1 we can easily determine which objects are shiny by the presence of areas of spectral reflectance.

Surface sheen, shine, gleam, etc. (see [Wyszecki 75] for a discussion of these terms) is a very important aspect in material discrimination. We regard metals as having a shiny appearance, whereas plastics, while they may have as smooth a finish, appear somewhat dull in comparison. Other surfaces may be altogether matte. These differences are caused by the presence, or absence, of local mirror-like or specular regions of reflected light, henceforth called highlights. If we can detect these highlights within an image, we can glean information that will help us to identify materials.

Some other consequences of finding highlights are that; (i) they would aid in constraining the size, color, and location of a light source; (ii) they would simplify object recognition or matching by identifying the regions so that some of the effects of the illumination could be 'factored out' and (iii) they also would enable constraints to be placed on object size and location [Thrift 82]. Perhaps a more basic or fundamental reason for wanting to locate highlights is that computer vision is concerned with modeling human vision, of which an inherent feature is the ability to locate highlights.

The detection of highlight regions proceeds by examining the stimuli that creates the sensation of a highlight. Horn's [Horn 75] model of surface reflectivity describes the two basic reflection components from a surface, specular and diffuse, as being separate quantities. Forbus [Forbus 77] used this information to generate a series of one dimensional profiles of intensity for curved surfaces, to see what parameters are relevant to the perception of highlights in archomatic images. Forbus noted that both the specular and diffuse reflection components must be present to create the sensation of a high light. Using this information, in conjunction with the expected sinusoidal shape of the specular reflection

component as given in Horn's equation, we can detect highlights.

1. DIFFERENTIAL OPERATORS

The sinusoidal shape of the specular component of reflection can be used to locate highlights by looking for zero-crossings in the first differential of the intensity image (see figure 1-1). Since zero-crossings in the first differential can correspond to either maxima or minima in the original image, care needs to be taken that we accept only those zero-crossings for which we have a maxima in the original image. Consequently, a concavity check is made by taking the second differential of the intensity image to make sure we have a peak of intensity and not a trough.

By using the first derivative of the Gaussian as our operator, we incorporate both a smoothing and differential operator into one step. The form of our operator is a two dimensional mask of values, calculated to simulate the derivative of the Gaussian, which can then be convolved with the intensity image. Since our convolution operator is now a function of σ , we can vary the sensitivity by using larger or small values for σ .

While we may choose to apply our differential operator using a single value for σ , that would be inappropriate considering the wide range of highlight sizes that may occur in an image. Highlights with a wide range of sizes can only be located by using multiple values for σ and then ORing all the results together. To ensure that we define zero-crossings not caused by singularities due to the choice of σ , the algorithm incorporates a phase which finds the zero-crossings for two values of σ and then ANDs those images together. Taking two values of σ relatively close together ensures that when we AND the results we do not eliminate 'true' highlights due to the scale of our operator. We can then OR a few such σ pairings to cover the gamut of highlight sizes.

However, we must remember that the differential mask is directional (non-isotropic) and must be applied to the image oriented at various angles of θ . It turns out that the convolution of the differential mask with the image need only be done twice, for two orthogonal angles, since the other angular orientations can be derived from those results. Complete zero-crossing angular orientation information can be ascertained from four angles of θ , each forty-five degrees apart. This means that two subsequent calculations need to be performed on the results of the orthogonal angle convolutions. Once we have the zero-crossing information for our differential masks the results can be combined with a concavity check to ensure we keep only those zero-crossings for peaks of intensity.

Concavity is determined by taking the second derivative with respect to a curve and looking at the sign of the resultant. In our case, this means convolving the original image with a two dimensional, second differential, mask. The mask is formed in the same manner as the first dif-

ferential directional mask, that is the differentiation is with respect to the Gaussian, but with the exception of using the Laplacian rather than the directional derivative so that the convolution need only be done once [Marr 75]. The result of the convolution is a two dimensional array composed of positive and negative regions demarcating concave up and concave down regions respectively. Only those zero-crossings within negative valued regions are accepted for further analysis. The result is an image containing zero-crossings for the original monochromatic image. There are still two other monochromatic images left from the original RGB images to process. Color is used to corroborate the data [Kanade 81] so that we have a 'true' highlight identification system. The complete zero-crossing information for the RGB image is shown in picture 2. Since zero-crossings chains are only one pixel wide, it is necessary to incorporate a region growing phase into the algorithm to locate the whole highlight.

2. REGION GROWING

The highlight pixels that surround the previously located highlight center chain can be found by growing outwards, from that central chain, in a locally orthogonal direction. To ascertain which direction is orthogonal the local direction of the highlight center chain is found over a three-by-three mask. The growing proceeds while the results of convolving a growing one dimensional orthogonal mask, with the original image, are increasing.

By convolving the power-of-two mask in figure 2-1 with the binary valued highlight chain image (1 for highlight pixels and 0 otherwise) and examining the resulting numeric value, we can determine the local direction. Using the eight-connected neighbor model there are four directions; north-south, east-west, southwest-northeast and northwest-southeast. For example if the convolution result is 17 or 68, the highlight center chain pixel is tagged with a east-west or north-south flag respectively. Similarly other values demarcate various compass directions. From the local direction labelling we can also label each of the pixels with a corresponding orthogonal direction. Choosing the local direction can be impossible for some patterns that can arise in a three-by-three area, so these pixels receive special treatment.

These pixels are labeled 'blob' pixels since they are without definite direction. To enable processing, they are marked with a flag that states every direction is orthogonal, and thus they are processed for each of the four directions. Once all the highlight center chain pixels have been tagged for direction, the orthogonal direction highlight growing can proceed.

Orthogonal to the center chain pixels which identifies the peak of the highlight the highlight intensity values decrease until reaching the value of the diffuse intensity for that surface. Therefore, a one dimensional mask, oriented in the orthogonal direction, is set up to calculate the difference between a highlight center chain pixel and points on each side of it (each side is treated separately

due to the non-symmetrical shape of most highlights). At each iteration of the algorithm the size of the mask is increased by one as the edge of the mask is moved 'outwards' by one pixel.

The algorithm keeps iterating for each highlight pixel point as long as the calculated value is increasing. When the algorithm terminates, the current value for the masksize determines the diameter of the highlight for the current side it was working on. Since blob pixels are processed for each direction, the current value of the mask size is stored so that it can be compared to the subsequent mask size values. The final mask size chosen for blob pixels will be the minimum of the directional mask sizes. The blob pixel will then be marked with one of the four directions that corresponds to the minimum sized mask (see picture 3 for a scene with the highlight diameters used to generate disks around the highlight center pixels). We now have each highlight chain pixel tagged with a highlight diameter to each side of it and a direction. The highlight is thus completely determined.

3. CONCLUSION

An algorithm is described in this paper which locally processes scenes of various objects to determine areas of spectral reflectance or highlights. The algorithm is based on the separability of the spectral from the diffuse reflectances by differential methods. Once the highlight center chains are detected by the algorithm, they are expanded to highlight regions by region growing in a direction orthogonal to the local orientation of the highlight. At the conclusion of the algorithm the information known about each highlight includes location, size and direction. Thus the algorithm provides information that a Computer Vision system must make use of when analyzing, or understanding, a scene. In addition, the information this algorithm provides can be used as a preprocessor to image processing algorithms that rely on predetermined areas of spectral reflectance. It can also be used to identify areas of spectral reflectance so that they can be removed from the scene. This eliminates illumination peculiarities which might confuse later/other algorithms that do pattern matching.

The fact that the algorithm is successful using moderately unconstrained images is important since it decreases the gap between the world that the computer can now understand and the extremely complicated one in which we live.

ACKNOWLEDGEMENTS

I would like to acknowledge the support and assistance of my supervisor, Dr. Brian Funt and the helpful criticisms of Kim Adamson-Sharpe. Nedenia Holm digitized the images contained in this paper and many others. This research was funded by Natural Sciences and Engineering Research Grant A4322.

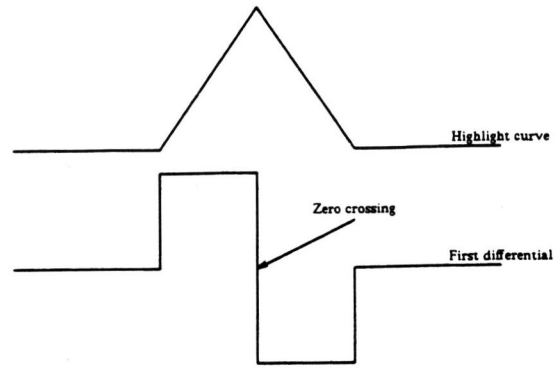




Figure 1-1: Highlight type curve and its derivative

2	4	8
1	0	16
128	64	32


Figure 2-1: Power-of-two mask



Picture 1: Image containing objects with highlights



Picture 2: The highlight image (highlights in red)



Picture 3: Highlight diameters used to generate disks (disks are shown in blue)

REFERENCES

- [Babu 85] Babu, M.D.R., Lee, Chia-Hoang, Rosenfeld, A.
Determining Plane Orientation from Specular Reflectance.
Pattern Recognition (1), May, 1985.
- [Barrow 81] Barrow, H.G., Tenenbaum, J.M.
Artificial Intelligence 17
North Holland, 1981.
- [Beck 72] Beck, J.
Surface Color Perception
Cornell University Press, 1972.
- [Bryant 83a] Burt, P.J.
Computer Graphics and Image Processing 21
Academic Press, 1983.
- [Burt 83b] Burt, P.J.
The Laplacian Pyramid as a Compact Image Code
IEE Transactions on Communications 3(4), April 1983.
- [Cook 82] Cook, R., Torrance, K.
A Reflectance Model for Computer Graphics.
ACM Transactions on Graphics 1(1):7-24, January, 1982.
- [Davies 81] Davies, E.R., Plummer, A.P.N.
Thinning Algorithms: A Critique and a New Methodology.
Pattern Recognition 14(1-6), 1981.
- [Forbus 77] Forbus, K.
Light Source Effects.
MIT AI Lab Memo (422), May, 1977.
- [Hall 83] Hall, R. Greenberg, D.
A Testbed for Realistic Image Synthesis
IEEE CG & A, November, 1983.
- [Horn 75] Horn, B.K.P.
Image Intensity Understanding.
MIT AL Lab Memo, August, 1975.
- [Kanade 81] Kanade, T.
Artificial Intelligence 17.
North Holland, 1981.

- [Marr 75] Marr, D.
Vision
W.H. Freeman and Company, 1982.
- [Pellicano 85] Pellicano, P.N.E.
Detection of Specularities in Color
Images Using Local Methods
Master's Thesis, Simon Fraser Uni-
versity, December, 1985.
- [Phong 75] Phong, B.T.
Illumination of Computer Generated
Images.
Commun. ACM, June, 1975.
- [Powers 79] Powers D.L.
Boundary Value Problems.
Academic Press, 1979.
- [Pratt 78] Pratt W.K.
Digital Image Processing
John Wiley and Sons, 1978.
- [Prewitt 70] Prewitt, J.M.S.
Object Enhancement and Extraction.
Picture Processing and Psychopic-
torics.
Academic Press, 1970.
- [Sears 58] Sears, F.W.
Optics
Addison-Wesley Publishing Com-
pany, 1958.
- [Shafer 84] Shafer, S.A.
Using Color to Separate Reflection
Components.
Carnegie-Mellon Technical Report
(136), April, 1984.
- [Thrift 82] Thrift, P., Chia-Hoang Lee.
Using Highlights to Constrain Object
Size and Location.
IEEE Publication, 1982
- [Ullman 75] Ullman, S.
On Visual Detection of Light
Sources.
MIT AI Lab Memo (333), May, 1975.
- [Wyszecki 75] Wyszecki, J.
Color in Business, Science and In-
dustry.
John Wiley and Sons, 1975.