Object Reconstruction with Photometric Stereo

Otavio Schipper Luiz Velho
Instituto Nacional de Matemática Pura e Aplicada

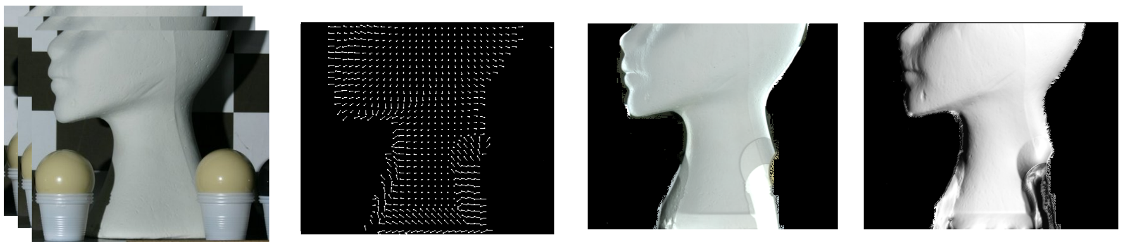

Figure 1: input images, calculated normal maps, albedo and 3d reconstruction.

Abstract

Photometric Stereo is a method for recovering a patch of surface from image data. The method involves reasoning about the image intensity values from several different images of a surface in a fixed view illuminated by different sources. This method directly measures surface orientation and recovers the distance to the surface at points corresponding to each pixel. We describe a system that is able to calibrate the lighting directions of the sources, find the best fit normal and albedo at each pixel and then to find a surface wich best matches the solved normals.

1. Introduction

The idea of Photometric Stereo is to vary the direction of incident illumination between successive images, while holding the viewing direction constant. The technique is photometric because it uses the radiance values recorded at a simple image point, in successive views, rather than the relative positions of displaced features. Traditional stereo techniques determine range by relating two images of an object viewed from different positions. In contrast, Photometric Stereo obtain surface orientation only from varying the direction of incident illumination between the successive images. Photometric Stereo is part of a larger set of methods known as shape-from-shading. These include methods whose outputs are surface normals and they can be used together with other reconstruction methods (such as a 3d scanner). For example, Normal Maps can be used to obtain high-quality renderings even when the actual mesh has low quality or low resolution.

Although the accuracy of 3d scanned data can be high, noise in measured positions can cause errors when surface normals are computed thus Normal Maps can also be used to help finding the best surface fitting using hybrid methods that combines depth information obtained via triangulation and measured surface orientations. [1]

The P.S. method has several advantages in comparsion to other reconstruction methods such as:

- Unlike single image shape from shading algorithms, Photometric Stereo makes no assumption of the smoothness of the surface [2].
- We can use low-cost equipaments such as digital cameras and flashes to obtain a good surface reconstruction .
- With a controled enviroment we can built a very simple set-up for acquiring geometry and reflectancew properties.

2. Shading Model

The principle of photometric stereo (first presented by Woodham in 1980 [3]) is to take three or more images of an object from the same camera view, each with lighting from a different, known direction. Fixing the camera and the object in position and illuminating the surface of the object using a point source that is far away (compared to the object dimensions) we can adopt a local shading model. We assume that there is no ambient illumination, so that the radiosity at a point P on the surface is given by the folowing expression [4]:

$$R(P) = \alpha(P)\vec{N}(P) \cdot \vec{S} \tag{1}$$

where $\alpha(P)$ is the surface albedo (wich we should assume to be constant), $\vec{N}(P)$ the unit surface normal at point P and $\vec{S}$ is the lighting direction (a unit vector).

We should note that we are assuming that the response of the camera is linear to the surface radiosity (in the interval we are exposing), so that the value of a pixel intensity is given by:

$$I(x, y) = kR(x, y) = k\alpha\vec{N}(x, y) \cdot \vec{S} \tag{2}$$

where k is the constant connecting the camera response and the input radiance.

We are also considering that the light sources are balanced so they all appear to have the same brightness (wich we are assuming to be equal to 1 in some arbitrary units. In fact we are using the same camera flash model at different positions varying the direction of the incident light.

3. Implementation

Our experimental setup consists of a Cannon Rebel digtal camera, a Manfrotto tripod, a vivitar 285 HV flash with an off-camera cable (in order to use it independent of the camera fixed position) and a black specular sphere used to probe the light directions [5]. We place the camera in a fixed position and the object in a controled environment with the black sphere close to the object. We take three or more images of the object varying the lighting direction changing the position and direction of the flash and use them as input data.

Our method for determining the lighting direction is done by the user marking the rectangular area that contains the sphere on the GUI . By knowing the shape of the sphere we can compute the normal at any given point on its surface, and therefore we can also compute the reflection direction for the brightest spot on the sphere. Assuming a single color channel with three or more image samples under different lightings we can obtain the surface normals for each point by solving a linear least squares problem with the following objective function where $\vec{G} = k\alpha\vec{N}$:

$$Q = \sum[(I_i(x, y) - S_i G^t)]^2 \tag{3}$$

Once we have the normals $\vec{N}$ for each pixel we can solve for the albedos by another least squarse solution with the following objective function:

$$Q = \sum[(I_i - \alpha S_i H^t)]^2 \tag{4}$$

where $\vec{H} = k\vec{N}$. To minimize it we should differentiate with respect to α and set to zero what implies that:

$$\alpha = \frac{\sum I_i S_i H^t}{\sum (S_i H^t)^2} \tag{5}$$

We can do this procedure for each channel independently to obtain a per-channel albedo (i.e. the object texture color). If a surface having these normals exist we can use path integration to find its depth values [6]. As we can not assume that, we will try to pose this problem as a least squares optimization. As the normals are perpendicular to the surface, they will be perpendicular to any vector on the surface as well. Let us consider a pixel (i, j) and its neighbour to the right (i+1,j). We can construct a vector $\vec{V}$ that is perpendicular to the normal $\vec{N}$ at the point (i, j) :

$$\begin{aligned} \vec{V} &= (i + 1, j, z(i + 1, j)) - (i, j, z(i, j)) \\ &= (1, 0, z(i + 1, j) - z(i, j)) \end{aligned} \tag{6}$$

As we know that$\vec{V} \cdot \vec{N} = 0$ it implies that:

$$n_x + n_z(z(i + 1, j) - z(i, j)) = 0 \tag{7}$$

and similary in the vertical direction:

$$n_y + n_z(z(i, j + 1) - z(i, j)) = 0 \tag{8}$$

It is possible to construct similar constraints for all of the pixels (wich have neighbours) and to form a matrix equation $M\vec{z} = \vec{v}$. We could use another least squares solution for this problem but we would have to deal with a very large matrix M^tM with as many rows and columns as pixels in the images. As most of the entries in the matrix M are zero we use a method for solving a linear system with sparse matrices to find the depths values.

4. Final Remarks

We have developed a system to acquire high-resolution geometry and reflectance properties using photometric stereo. One example of our results can be seen in figure 1.

References

[1] Nehab Diego et. al ”Efficiently Combining Positions and Normals for Precise 3D Geometry”, ACM Transactions on Graphics - SIGGRAPH , July 2005.

[2] Zhang R. et. al ”Shape from Shading: A Survey”, Computer Vision Lab, UCF, IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume 21 , Issue 8, 1999.

[3] Woodham, R. ”Photometric method for determining surface orientation from multiple images”, Optical Engineering, Vol 19, No 01.

[4] Rushmeier, H ”Computing Consistent Normals and Colors from Photometric Data”, IBM Reserch Center, 2nd International Conference on 3D Digital Imaging and Modeling (3DIM ’99),Ottawa, CANADA.

[5] Debevec P. et. al ”Photometric Stereo for Archeological Inscriptions”, USC - Institute for Creative Technologies, Siggraph Technical Sketches, 2004.

[6] Forsyth D. and Ponce, J. ”Computer Vision; A Modern Aproach”, Prentice Hall, 2003.