Obtaining the Shape of a Moving Object with a Specular Surface

Atsuto Maki
atsuto.maki@crl.toshiba.co.uk
Toshiba Research Europe Ltd, Cambridge

Roberto Cipolla
cipolla@cam.ac.uk
Department of Engineering, University of Cambridge

Recovering the 3D surface of an object from shaded images has been long investigated in computer vision. A technique of using the shading variance due to object motion has been recently studied and attracts attention [3, 4, 8]. In this paper, we also concern ourselves with this technique because it allows us to stay with the simplest possible setup.

The basic scheme of the technique is to employ the framework of multi-view stereo, which first requires correspondence of several sample points to recover the extrinsic camera parameters, and then to exploit the shading variance due to object pose for dense matching. Notwithstanding the recent advances, one of the limitations of the previous techniques is concerned with the object surface, which is commonly assumed to have Lambertian reflectance. That is, if the surface takes on specularities it will be a problem for the first stage to accurately compute point correspondence. At the second stage, in which the linear subspace constraint [7] is imposed, specular reflections will be even more problematic as they produce non-linear highlights [6].

The motivation of this paper is to overcome such a limitation at each stage in the scheme. This paper is most relevant to the work of [3] and [4] in that we also investigate the linear subspace constraint. However, it differs from them in the challenge that we propose techniques to deal with specular components both in feature correspondence and dense depth search. Moreover, as a direct application we show that we can acquire a linear image basis of target object by using the consequence of dense matching [5] and removal of specular components.

Our assumption is that the target object is illuminated by a distant unknown light source. The surface may or may not be textured, and is approximated by a Lambertian component plus a specular component but without shadow.

Point correspondence We first compute sparse feature correspondence across the input images which we use for estimating light source parameters as well as external camera parameters. In order to eliminate the spurious matching due to specularities at this stage, we compute outliers by minimising the LMedS error in terms of photometric measures and exclude them before computing the camera parameters and light source parameters.

Dense depth search We propose to recover the shape of an object by a dense depth search while evaluating the correct matching with the linear subspace constraint as investigated in [3, 4]. In order to deal with specularities, in our case, we choose to employ a minimum of five images as inputs. This is because we first need three images, if no specularities are present, to estimate the surface normal (up to an ambiguity) at a hypothesised surface location as well as the fourth image to verify that the surface indeed passes through that point. The fifth image is to detect the existence of a specularity by checking the consistency with the other four inputs.

This detection of specularity is inspired by the principle of 4-source photometric stereo [1, 2, 9]. That is, among five or more input images we interchangeably utilise an optimal *intensity subset* consisting of four projected image intensities to find the best matches. The strategy is motivated by the fact that the specularities travel on the surface thanks to the object motion and therefore on the assumption that they are non-overlapping.

Figure 1 and Figure 2 illustrate the input images and recovered surface. Since a light source is placed beside the camera, in each input image, parts of the surface where the orientation is close to the direction toward the camera tend to take on specular reflection. It also includes resulting depth map computed for the reference input images in the midst. The surfaces are well recovered considering the fact that only five images are used, and uninfluenced by the partial specularities. The results on the upper part of the water pitcher is erroneous because that part is not visible throughout all the input images and is also with strong inter-reflections. Figure 3 shows how an input image can be separated into specular components and a linearised image within the process of dense matching. (In the paper we also show how the other input images are separated.)

In the paper we formulate the problem of recovering the surface of a moving object using the shading variance due to object motion. Our key idea for dense point matching is to use an optimal subset of intensities that is the least influenced by specular reflections. To our knowledge it is one of the few trials of finding correspondence in the presence of specularities on the top of intensity variance due to object motion. We have also shown that it is possible to linearise the input images by directly applying our results.

Figure 1: Three examples out of five input images of a water pitcher captured under a point light source whose direction is unknown.

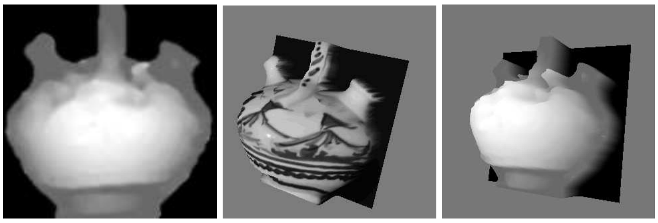

Figure 2: A depth map estimated referring the input image in the midst (the lighter, the closer) and the recovered surface with and without texture.

Figure 3: Left: Extracted specular components in the reference input image (residuals, stressed by factor 5 for display reason). Right: Linearised images after removing the specular components.

[1] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE-PAMI*, 25(10):1239–1252, 2003.

[2] E.N. Coleman Jr. and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *Computer Graphics And Image Processing*, 18:309–328, 1982.

[3] N. Joshi and D.J. Kriegman. Shape from varying illumination and viewpoint. In *11th ICCV*, 2007.

[4] A. Maki, M. Watanabe, and C.S. Wiles. Geotensity: Combining motion and lighting for 3d surface reconstruction. *IJCV*, 48:2:75–90, 2002.

[5] A. Nakashima, A. Maki, and K. Fukui. Constructing illumination image basis from object motion. In *7th ECCV*, pages III:195–209, 2002.

[6] M. Oren and S. Nayar. A theory of specular surface geometry. In *5th ICCV*, pages 740–747, 1995.

[7] A. Shashua. *Geometry and photometry in 3D visual recognition*. PhD thesis, Dept. Brain and Cognitive Science, MIT, 1992.

[8] D. Simakov, D. Srolova, and R. Basri. Dense shape reconstruction of a moving object under arbitrary, unknown lighting. In *9th ICCV*, pages 1202–1209, 2003.

[9] F. Solomon and K. Ikeuchi. Extracting the shape and roughness of specular lobe objects using four light photometric stereo. *IEEE-PAMI*, 18(4):449–454, 1996.