# Specularity and Shadow Interpolation via Robust Polynomial Texture Maps


Mark S. Drew, Nasim Hajari
{mark,nha16}@cs.sfu.ca

School of Computing Science
Simon Fraser University, Vancouver, B.C., Canada

Yacov Hel-Or
toky@idc.ac.il

Department of Computer Science
The Interdisciplinary Center, Herzliya, Israel

Tom Malzbender
tom.malzbender@hpl.hp.com

Media and Mobile Systems Lab
Hewlett-Packard Laboratories, Palo Alto, CA


Polynomial Texture Maps (PTM) [1] form an alternative method for apprehending surface colour and albedo that extends a simple model of image formation from the Lambertian variant of Photometric Stereo (PST) to more general reflectances. Here we consider solving such a model in a robust version, not to date attempted for PTM. But the main upshot of utilizing robust regression is in the identification of both shadows and specularities automatically, without the need for any thresholds, in a tripartite set of weights for pixels that are labelled as matte, shadow, or specularity. Original images are captured using a hemispherical set of lights, and pixel values across the lighting directions are then labelled as inliers, or outliers of two types. A per-pixel robust regression on luminance is carried out using Least Median of Squares, and automatically-identified outlier pixels are labelled as shadows if they are darker than matte and correspondingly, specular outliers are too bright. Inlier identification generates correct values for chromaticity and for surface albedo and thus matte luminance and colour. Then a robust version of PST, using only PTM inliers, improves estimates of normal vectors and albedo recovered. With specular pixel values over the lights in hand we model specularity using a radial basis function (RBF) regression, and non-specular pixel departures from matte using a second RBF set. Then for a new lighting direction, we can readily interpolate both specular content as well as shadows.

Here, we are interested in using PTM as a vehicle to carry out *interpolation* of specularities and shadows. To the best of our knowledge robust methods have not to date been applied to PTM, and we use these to be able to accomplish interpolation. As well as using robust regression, this paper moreover shows how outliers can be classified as belonging to two types: either specular highlights, or self- or cast-shadows. Knowledge of inlier pixel values means that recovered surface albedo and chromaticity is robust, in the sense of ignoring outlier contributions and thus more accurately mapping surface reflectance and colour.

Finally, knowledge of outlier labels means that we can independently model specularity and shadow. Then for a new lighting direction, we can generate pixel values interpolating known values of both; here we use a Radial Basis Function (RBF) interpolation model. In this paper we carry out robust regression on the luminance values, not on R,G,B separately. Since we generate the specularity $\varsigma$ in an interpolated image, separately from the remaining contribution $\sigma$, we can then produce a full-colour interpolant image using the luminance times matte chromaticity for the non-specular contribution, plus specular-luminance times the chromaticity for the specular colour.

Results on re-creating the *input* images are shown to have excellent agreement with the originals, over a variety of input sequences. For shadow and specularity interpolation – generating images for non- observed lighting directions – the method is shown to indeed generate sensible results. The main contributions in this paper are (1) application of robust regression to PTM; (2) separate modelling and thus better capturing of shadows and specularities; and (3) specular and shadow interpolation.

Suppose we have acquired $n$ images of a scene, taken from $I = 1..n$ different lighting directions $\boldsymbol{a}^I$ as in Fig. 1(a). Let each RGB image acquired be denoted $\boldsymbol{\rho}^I$. Suppose we make use of the luminance images instead, $L^I = \sum_{k=1}^{3} \rho_k$. This reduction in dimensionality reduces the computational burden of robust regression and, since we mean to separate out the specularity from the matte component, we can re-insert colour later separately for matte colour and specular colour.

Then a PTM model consists of a nonlinear regression from lighting to luminance via a vector of polynomial terms $\boldsymbol{p}$, with $\boldsymbol{p}$ a function of lighting direction $\boldsymbol{a}$, as follows:

$$\begin{bmatrix} \boldsymbol{p}(\boldsymbol{a}^1) \\ \boldsymbol{p}(\boldsymbol{a}^2) \\ \dots \\ \boldsymbol{p}(\boldsymbol{a}^n) \end{bmatrix} \boldsymbol{c} = \begin{bmatrix} L^1 \\ L^2 \\ \dots \\ L^n \end{bmatrix} \quad \text{, or} \quad P\,\boldsymbol{c} = \boldsymbol{L} \tag{1}$$

where $\boldsymbol{c}$ is a vector of regression coefficients. Each pixel has its own $\boldsymbol{c}$, and the $\boldsymbol{L}$ vector is the collection of all luminances at that pixel over the $n$ images, for polynomials $\boldsymbol{P}$ for the known lighting directions. For a robust regression to replace (1), here we utilize the LMS [2], which generates outliers "on a silver platter", without any intervention. So in this luminance-based variant, regression is

$$\boldsymbol{c} = LMS(\boldsymbol{P}, \boldsymbol{L}) \tag{2}$$

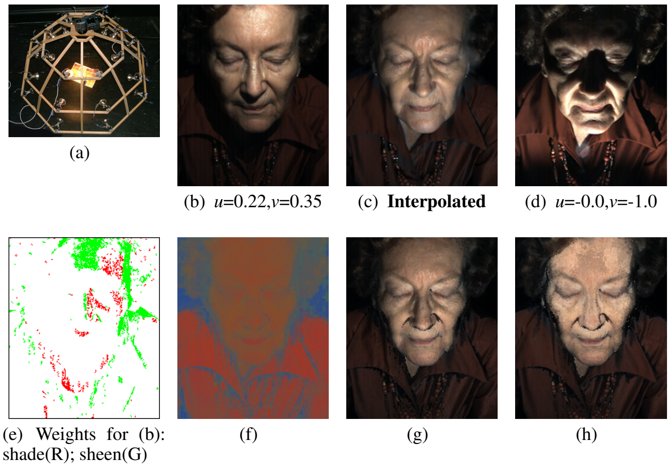


Figure 1: (a): Hemispherical dome with multiple, identical lights. (b,d): Two inputs. **(c): Interpolant for light between (b) and (d).** (e): Weights. (f): Recovered chromaticity $\boldsymbol{\chi}$ for (c). (g): Matte from PTM $L = \boldsymbol{p}(\boldsymbol{a})\boldsymbol{c}$. (h): Intrinsic, $\boldsymbol{\rho}_{intrinsic} = \alpha\boldsymbol{\chi}$, where $\alpha$ is albedo.

The output of LMS is the set of regression coefficients $\boldsymbol{c}$, plus a set of $n$ weights (labels) $\boldsymbol{w}$ identifying inliers ($w = 1$) and outliers ($w = 0$), thereby excluding some lights at this pixel. But we know here that outliers are generated because (1) $L$ values are too high to suit the model (1) – we take these to be specular contributions; or (2) luminance is too low (or the model generates negative values), for a particular light – these are likely shadow locations. Thus we arrive at a tripartite set of weights $\{w^0, w^+, w^-\}$ at each pixel, with $w^0$ set for lights generating inlier values, $w^+$ for specularities, and $w^-$ for shadows. Fig. 1(e) shows weights corresponding to input image Fig. 1(b), $w^0$ as white, $w^+$ as green, and $w^-$ as red.

Now we can go on to generate a matte version of the input set of images, or indeed an interpolated matte image, using the regression result $\boldsymbol{c}$, as in Fig. 1(g). We are assured of doing better than standard PTM since we have excluded distracting specularities and shadows from consideration whilst generating coefficients $\boldsymbol{c}$.

At this point, we already have an advantage of applying a robust method to the problem at hand, *viz.* a more reliable calculation of coefficient $\boldsymbol{c}$. But in fact we also have produced a better grasp of colour, as well. Let us factor each RGB triple $\boldsymbol{\rho}$ into luminance $L = R + G + B$ times chromaticity $\boldsymbol{\chi}$. Luminance will be composed of a scalar albedo $\alpha$ times lighting strength times shading factor $s$; since we have no way of disentangling lighting intensity from surface reflectance, we shall simply lump both scalars into $\alpha$. Thus,

$$\boldsymbol{\rho} = s\alpha\boldsymbol{\chi}, \quad \boldsymbol{\chi} \equiv \{R, G, B\}/(R + G + B) \tag{3}$$

An intrinsic image (for this lighting strength), i.e. surface independent of shading, would then be

$$\boldsymbol{\rho}_{intrinsic} = \alpha\boldsymbol{\chi} \tag{4}$$

I.e., this is what the surface would look like under this light, with shading removed.


[1] T. Malzbender, D. Gelb, and H. Wolters. Polynomial texture maps. In *SIGGRAPH 2001*, pages 519–528, 2001.

[2] P. J. Rousseeuw. Least median of squares regression. *J. Amer. Stat. Assoc.*, 798:871–880, 1984.