

A Photometric Sampling Strategy for Reflectance Characterization and Transference

Mario Castelán, Elier Cruz-Pérez, Luz Abril Torres-Méndez

Centro de Investigación y Estudios Avanzados del IPN, Ramos Arizpe,
Coah., Mexico

mario.castelan@cinvestav.edu.mx

Abstract. Rendering 3D models with real world reflectance properties is an open research problem with significant applications in the field of computer graphics and image understanding. In this paper, our interest is in the characterization and transference of appearance from a source object onto a target 3D shape. To this end, a three-step strategy is proposed. In the first step, reflectance is sampled by rotating a light source in concentric circles around the source object. Singular value decomposition is then used for describing, in a pixel-wise manner, appearance features such as color, texture, and specular regions. The second step introduces a Markov random field transference method based on surface normal correspondence between the source object and a synthetic sphere. The aim of this step is to generate a sphere whose appearance emulates that of the source material. In the third step, final transference of properties is performed from the surface normals of the generated sphere to the surface normals of the target 3D model. Experimental evaluation validates the suitability of the proposed strategy for transferring appearance from a variety of materials between diverse shapes.

Keywords. Reflectance transference, singular value decomposition, random Markov fields.

1 Introduction

The interaction of surface with light is an active topic of interest in the fields of computer vision and computer graphics, since recovering shape and appearance is possible through the comprehension of such interaction. In this sense, understanding the reflecting behavior of real world materials represents a challenge as they exhibit different properties such as specularity, diffuseness, anisotropy, and retro-reflection. Appearance

models are helpful in areas such as image interpretation, 3D reconstruction from images, perception, lighting interpolation, and image synthesis. Particularly, the photo-realistic rendering of 3D models has been the subject of study for many years, and it remains an open research topic. The patterns and characteristics to be recognized are focused on the appearance properties in terms of illumination changes together with camera motion. Generally, these problems can be solved by considering three elements: (1) a model of the properties to be determined, (2) the characteristics taken from the scene, and (3) an algorithm to find the parameters of the model that better adjust to these characteristics.

Among the variety of methods for characterizing appearance from imagery, some use mathematical models aimed at accurately describing the reflectance phenomenon [20, 4, 12], i.e., a parametric equation that helps rendering the appearance of 3D models in accordance with the expected reflectance. Unfortunately, these approaches depend on the choice of a specific parametric family and do not deal with spatially varying reflectance. Other strategies use special sampling devices which require spherical or planar samples of the material to be analyzed [16, 7, 17, 11], constraining the characterization of reflectance to the existence of the physical specimen. In general, attention has been paid to characterizing reflectance by acquiring a considerable amount of images through special devices which consider calibration of light sources, introducing additional issues such as computational storage and errors in calibration.

Alternatively, in this work we focus on the statistical analysis of the appearance of a pixel which, in the form of a compact matrix, contains

the behavior of a surface patch with a moving light source. Our main goal is to automatically estimate, from a source object, appearance models suitable for transferring properties to target objects with known 3D geometry. The appearance behavior is related to the optical properties of a material as well as properties of the illumination source which is incident on the material; these properties being recorded by a camera are considered in this paper as the appearance information of the scene.

The contributions of this article are as follows. First, a novel representation based on the statistical decomposition of luminance matrices is proposed for the purposes of reflectance characterization. Such representation can be conveniently coupled with a Markov Random Field strategy that solves for correspondence between the surface normals of the source object and those of a synthetic sphere, which is used as an intermediate reflectance object in order to diminish transference errors. The second contribution of this work is therefore a methodology for reflectance transference that considers, after alteration of their singular values, a weighted reconstruction of luminance matrices that reflects the correspondence error between surfaces normals. Fig. 1 illustrates the result of using our method for transferring the reflectance of styrofoam mannequins covered with different materials onto the surface normals of Stanford's Buddha, dragon, and rabbit.

1.1 Related Work

The idea of reflectance was mathematically introduced by the BRDF [10] (Bidirectional Reflectance Distribution Function), a function that relates incoming energy over a surface patch (radiance) with the outgoing energy (irradiance) captured by the camera. The aim of BRDFs is to analytically represent the physical reflectance of real world materials as a mathematical function that considers local shape and illumination parameters. Reflectance may also be understood as the fraction of luminous flow striking a surface which is reflected of it. The basic form of the BRDF can be found in the Lambertian model, which can be expressed by the following equation for each pixel at a position (u, v) :




Fig. 1. Reflectance transference. A source object is illuminated in concentric paths in order to sample and transfer its reflectance onto the surface normals of a target object. The figure illustrates the result of using our method for transferring the reflectance of styrofoam mannequins covered with different materials onto the surface normals of different objects of the Stanford database.

$$I(u, v) = \rho(u, v)\lambda \mathbf{l} \cdot \mathbf{n}, \quad (1)$$

where I is the image irradiance, ρ is the albedo (a scaling factor indicating the amount of light to be reflected), λ is the light source intensity, $\mathbf{l} = (l_x, l_y, l_z)$ is the illuminant vector, and $\mathbf{n} = (n_x, n_y, n_z)$ is the surface normal vector. Eq. 1 is a common reference for many BRDFs which consider more complex reflectance such as specular regions and spatially varying albedo [20, 4, 12].

Traditionally, the BRDF of a sample material is measured using goniorefectometers, special devices that position a light source and a detec-

tor (photometer) in several directions over a planar sample material. Unfortunately, these devices measure one reflectance value at a time, making the gathering process slow. In the last decade, however, digital cameras have been used to assist the sampling process, as in the work of Ward [25], where a hemispherical mirror was used with a fish-eye camera and a moving light source in order to measure reflectance in all incident angles and for BRDF modeling.

In [8], Dana et al. developed a device to measure spatially varying BRDFs. To this end, a digital camera, a robotic arm, and a light source were used in order to measure reflectance on planar surface sample materials. In their work, samples were taken from varying angles of the light source and camera poses, from which a database of 60 types of materials can be found in [6]. Later, Dana et al. [7] made their previous device more efficient using a parabolic mirror. Similarly, Ghosh et al. [11] proposed using illumination basis functions for speeding up the measuring process. Additionally, in the works of [16] and [17], sampling is performed using superficial curves in order to recover BRDFs and reduce the number of images to be sampled along the hemisphere. In these methods, spherical samples of the material to be characterized are required. The above approaches share some common features: (1) reflectance is measured from images in accordance with changes in both illumination direction and camera pose, (2) the surface normals of the samples are known for each surface patch of the sample, (3) irradiance is determined after performing a radiometric calibration process on the light sources, which may imply using additional optical devices and samples of a highly diffuse material like spectralon, and (4) they consider the existence of either planar or spherical samples of the analyzed material.

Alternatively, Mertens et al. [18] proposed a texture transference method from a known 3D geometry object with known texture onto a target 3D model. Their work is based on the correlation between local shape and texture of the source object. In this sense, this approach focuses on both geometric and texture information for transference, while ignoring the effect of changes in illumination

which avoids BRDF modeling. Although the reported results keep an accurate level of photorealism, no reflectance information is characterized in the generated rendered view which only considers the light source direction of the source object.

In this paper, we introduce a novel solution to the above problems with a method that only requires features (1) and (2), i.e. no calibration of the light source is needed nor planar or spherical sample materials are assumed available. Instead, our approach is based on a recent work developed by [13] that considers the statistical analysis of a pixel's reflectance along concentric circles of a light source. While this analysis is useful for the purposes of reflectance characterization, the transference method solves for surface normal correspondence via Markov Random Fields.

The rest of the paper is organized as follows: in Section 2, the reflectance characterization developed for a pixel-wise decomposition is described; later, in Section 3, a Markov Random Field model for surface normal correspondence is depicted in order to establish a reflectance transference method between a source object and a target geometry; Section 4 presents experimental evaluation for characterizing and transferring the reflectance of objects with varying textures; finally, Section 5 provides conclusive remarks and future extensions of this work.

2 Reflectance Characterization

Photometric sampling can be defined as a process where the surface normals of an object are estimated through the excitation of the object's surface while a light source rotates around it. The method can be regarded as a special case of photometric stereo where extensive sampling is performed using a moving light source. Traditional photometric sampling approaches [22, 15] consider only variations around the azimuth angle of the moving light source, i.e., a single circle is used to sample reflectance. Recently, an extended photometric sampling methodology was introduced in [13] for the purposes of correcting photometric databases. The main idea of the method is to include variations of the illumination source along the zenith direction

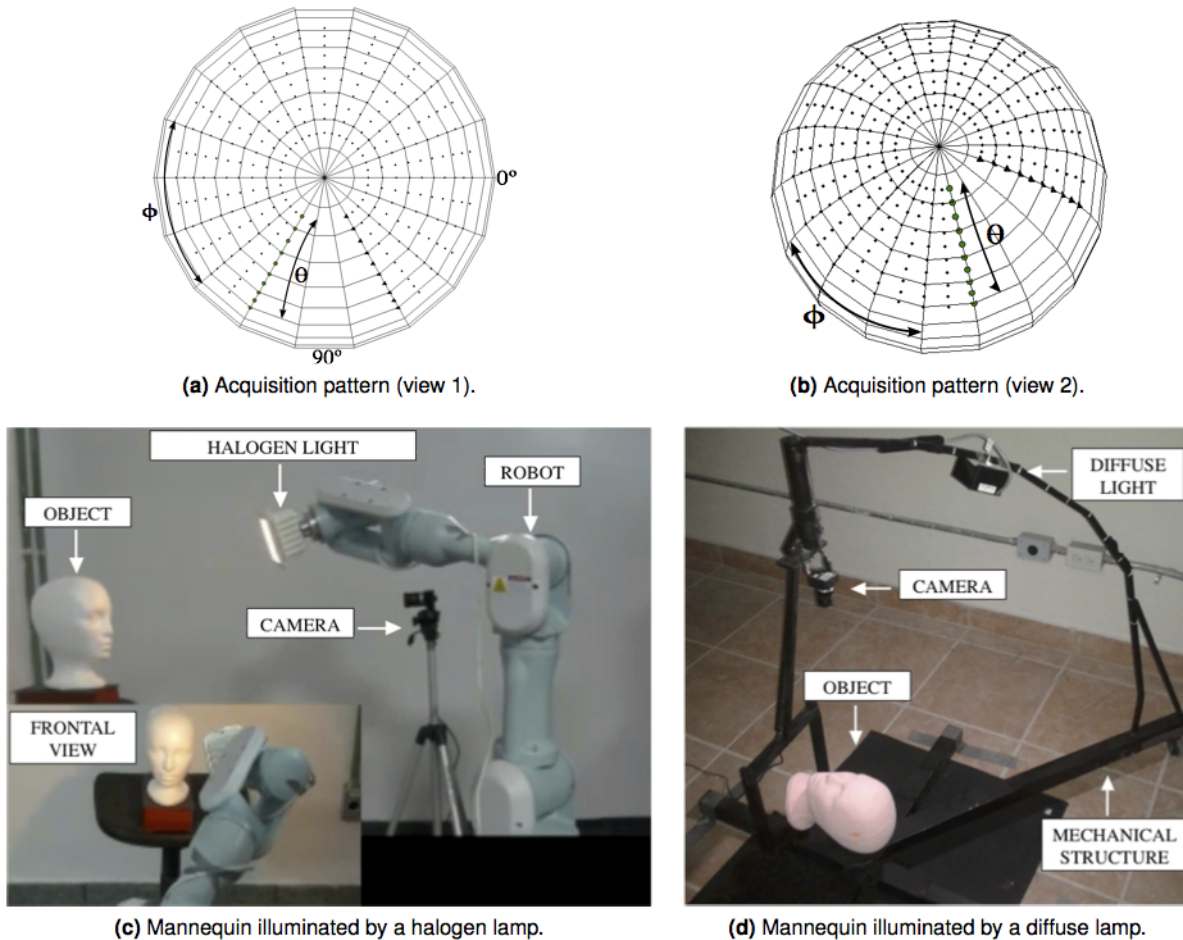


Fig. 2. Photometric sampling. Views of the illumination path for photometric sampling are provided in (a) and (b), where ϕ and θ stand for the azimuth and zenith angles of the light source direction, respectively. Views of the two database acquisition settings used for experimental evaluation are also shown: (c) a robotic arm manipulating a halogen light and (d) a metallic structure where a diffuse light was manually adjusted to the desired position

in order to build, for each image pixel, a *luminance matrix*. The authors showed how fitting sine functions on the singular vectors of the luminance matrix diminishes artifacts and noise inherent to the data acquisition process, therefore improving the photometric consistency of the database.

Other approaches such as [23] exploit the idea of determining an object's shape based on similarity of radiance changes observed at its surface points. In [14] the geometry of objects with general reflectance properties is computed from images with photometric changes. The method

is capable of handling objects with arbitrary and spatially-varying BRDFs while almost no calibration is required. Unlike the former techniques, our method is concerned with the characterization and transference of appearance rather than 3D shape recovery. To this end, we propose using the statistical decomposition of the luminance matrix for reflectance characterization and further transference. The main idea here is to study the role of the singular vectors and singular values of the luminance matrix in the generation of reflectance features of the object such as color, texture, and

shininess, as well as the illumination properties of the light source.

2.1 The Luminance Matrix

A luminance matrix may be thought of as a collection of pixel intensity measurements whose columns and rows respectively span variations along the azimuth and zenith angles of the light source. In this sense, a luminance matrix represents the energy captured in each pixel for the set of all the illuminations used during the photometric sampling process. In order to illustrate how the sampling is performed, Fig. 2 presents views of the concentric pattern used to illuminate the object. The variations in azimuth and zenith angles of the light source are represented by ϕ (from 0° to 360°) and θ (from 0° to 90°), respectively. The acquisition settings used for our experimental evaluation are also shown in the figure. The first setting is depicted in (c), where a robotic arm surrounds the object of interest, sampling the reflected luminance from a halogen light source attached to the end effector of the robot. The second setting is shown in (d), where a metallic structure is used to impose variations in azimuth and zenith angles of a diffuse light source.

Let us turn our attention to the core equation of the extended photometric sampling: for each image pixel of the source object, the singular value decomposition of the luminance matrix \mathbf{M} can be defined as

$$\mathbf{M}_{a \times z} = \mathbf{U}_{a \times r} \mathbf{\Sigma}_{r \times r} \mathbf{V}_{r \times z}^T, \quad (2)$$

where \mathbf{U} and \mathbf{V} are the matrices whose columns contain the left and right singular vectors of the rank- r matrix \mathbf{M} , and $\mathbf{\Sigma}$ is the diagonal matrix of the singular values of \mathbf{M} , namely $[\sigma_1, \sigma_2, \dots, \sigma_r]$. The numbers of lighting variations are represented by a and z respectively for the azimuth and zenith angles of the light source. An example of the luminance matrix for an image pixel of a mannequin is provided at the top panel of Fig. 3 (a). The highlighted pixel appears below the right eye of the mannequin. In the exemplified pixel, the luminance matrix presents variations for ten zenith angles of the light source, i.e., ten concentric circles were

performed during the photometric sampling process. Note how the measured intensities along the columns of the luminance matrix resemble sine functions, where departures from sine mean departures from Lambertian reflectance.

The main idea of the extended photometric sampling of [13] is to fit sine functions onto the singular vectors of the luminance matrix of all of the pixels so as to impose Lambertian-like reflectance. As a consequence, numerical artifacts that may affect the photometric consistency of the database are diminished. Fig. 3 (b) presents results obtained after applying the photometric correction on our databases. The correction is performed separately on each color channel on the RGB space. From left to right, four different substances covering a styrofoam mannequin are shown: matte paint, sawdust, nail varnish, and crayon. The top row of the panel presents the original images (before correction), the middle row depicts the corrected images, and the bottom row shows difference maps between the original and corrected images. Note how, as an effect of the sine fitting, the shadowed areas in the original images appear recovered while the specular region result lowered.

As far as the statistical decomposition of the luminance matrix is concerned, the first singular vectors for two luminance matrices are shown at the bottom-left diagram of Fig. 3 (a), for Lambertian (diffuse) and specular (Phong [20]) rendered spheres. Let us refer to \mathbf{M}_d and \mathbf{M}_s as the luminance matrices for a pixel of diffuse and specular synthetic spheres, respectively. The figure presents plots of the first left singular vector of the luminance matrices of a highlighted pixel for both spheres. Note how the singular vector of the specular pixel exhibits a peak revealing the nature of its specular reflectance. This suggests that the singular vectors contain information about how a surface patch of an object tends to reflect light. According to this idea, assigning the singular vectors of a luminance matrix from a source object onto a similar surface patch of a target object may result in transferring reflectance information. In this paper, we explore the role of singular values and singular vectors in the generation of reflectance and how this could be further used for the purposes of reflectance transference.

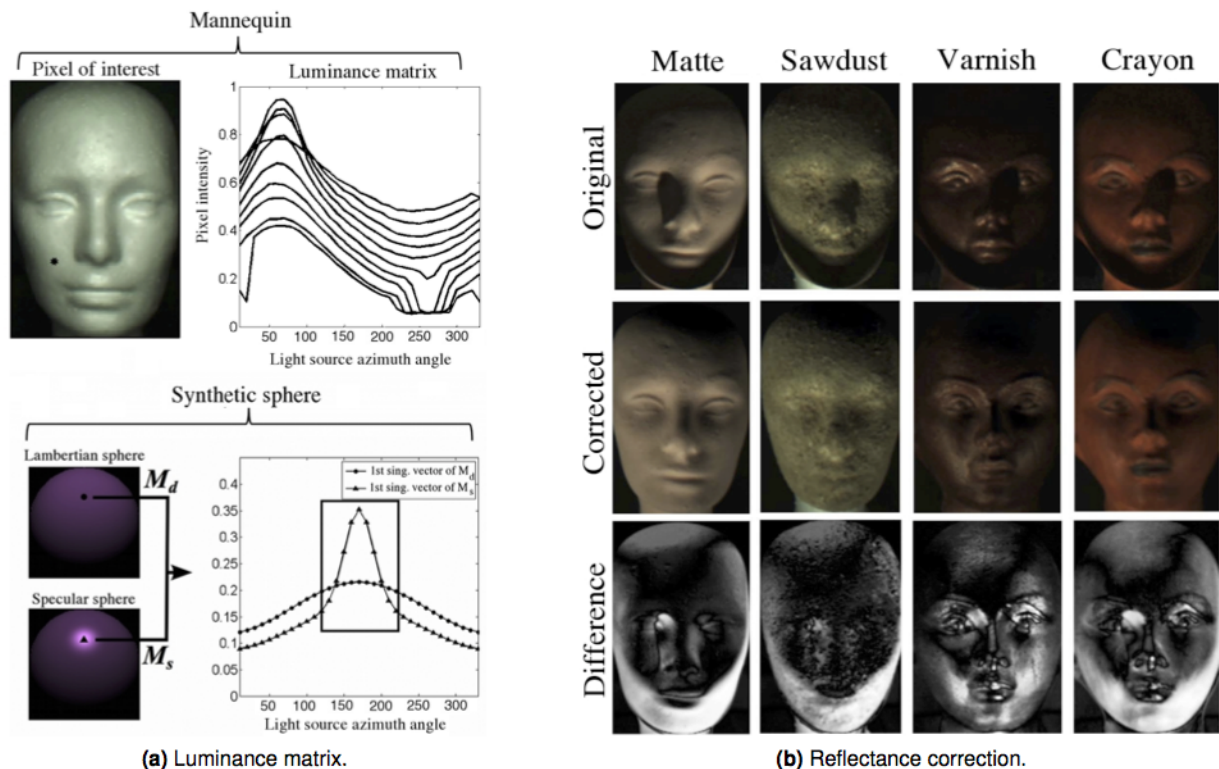


Fig. 3. Luminance matrix and reflectance correction. The figure presents in (a) an example of the luminance matrix of a pixel of a mannequin image. The first left singular vector of luminance matrices for a pixel of Lambertian and specular synthetic spheres are also shown in (a); M_d and M_s refer to the diffuse and specular luminance matrices of the spheres, respectively. The right diagram of the figure presents results after applying the reflectance correction of [13] on databases of four different materials: matte painting, sawdust, red nail varnish, and orange crayon. The difference between original and corrected images is shown at the bottom row of the panel

2.2 Pixel-Wise Factorization

It is important to recall that there is one luminance matrix attached to each image pixel. This matrix spans variations of the light source along its columns (azimuth angles) and rows (zenith angles). Unlike performing SVD on the database as a whole (i.e., taking each image in the database as a long vector), the aim of separately performing SVD on each luminance matrix is to isolate, pixel by pixel, the lighting variations imposed by the moving light source over the object's surface normals. Such separation avoids mixing up the reflectance contribution of all the pixels in the database, facilitating the characterization of reflectance at the local (surface patch) level. A pixel-wise decompo-

sition is also convenient for the reflectance transference method proposed in this work.

In order to provide further details on the luminance characterization at the local level, Fig. 4 explores the appearance of the singular values and singular vectors of the luminance matrices for the red varnish mannequin. The figure is organized in two panels: the top panel refers to variations along the zenith angle of the light source while the azimuth angle remains fixed; the opposite case (fixed zenith angle, varying azimuth angle) is depicted at the bottom panel of the figure. In the figure, the first singular value, first left and first right singular vectors of each image pixel are referred to as σ_1 , u_1 , and v_1^T , respectively. The first singular value is shown in (a), for each luminance matrix as a

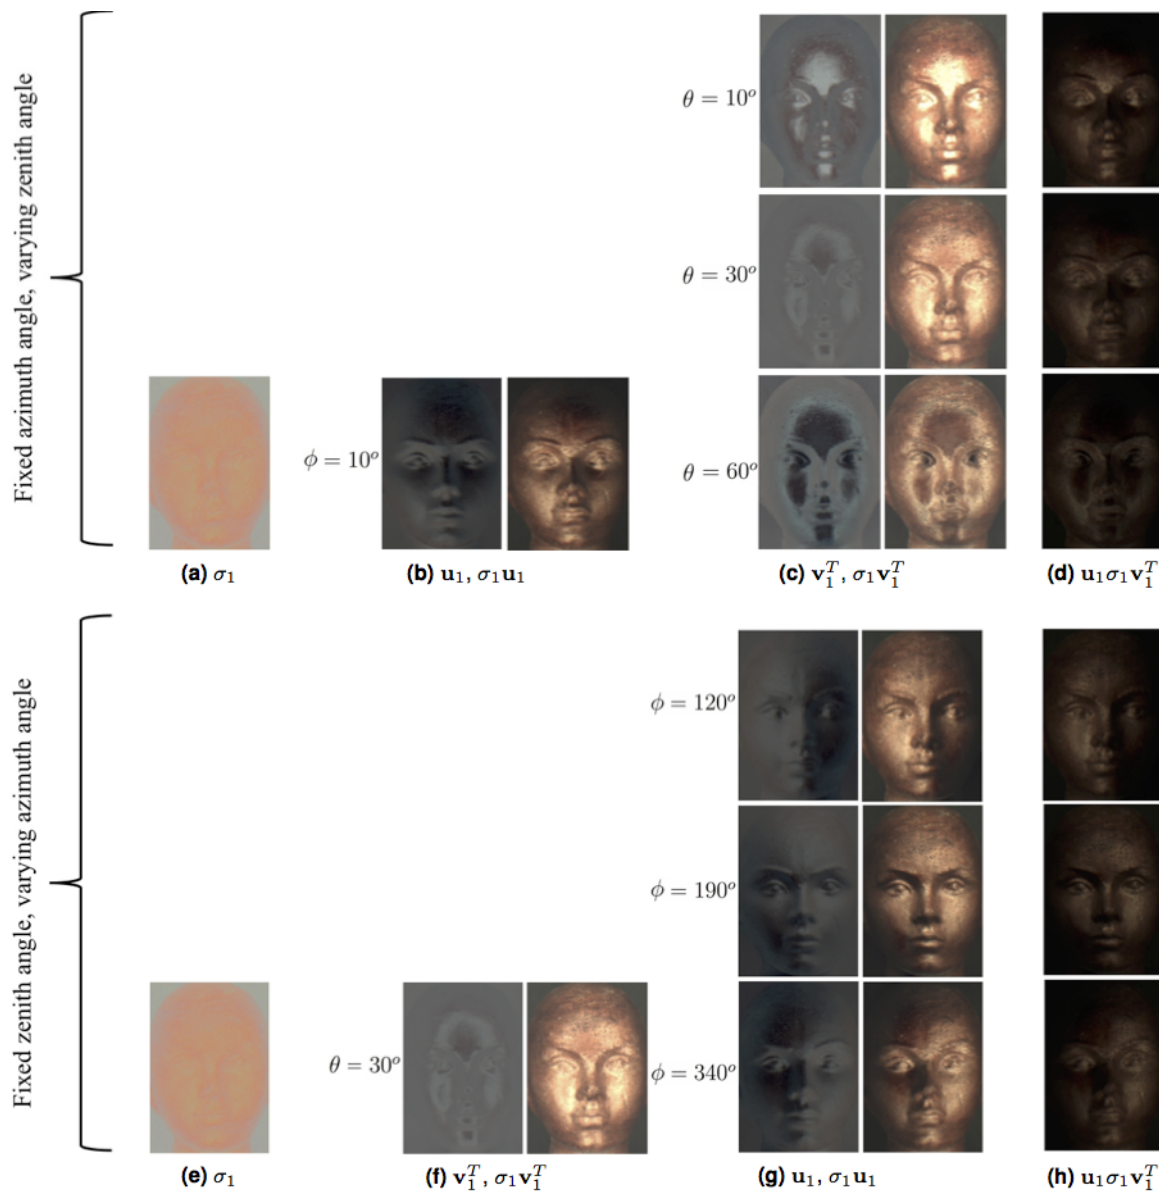


Fig. 4. Decomposing the luminance matrices of a styrofoam mannequin covered with red nail varnish. The figure illustrates the first singular values and the first singular vectors of the luminance matrices of the mannequin covered with red nail varnish. The top panel shows the results on varying the zenith angle of the light source while fixing the azimuth angle. The opposite case is shown in the bottom panel. The singular values are represented by σ_1 while the left and right singular vectors are referred to as \mathbf{u}_1 and \mathbf{v}_1^T , respectively

pixel map. The visual analysis of the resulting image suggests that singular values are weighting factors representing the color variation with which each surface normal reflects the light source used

in the sampling process. Particularly, the surface normals with greater slant, (the ones around the boundaries of the mannequin's face, and edges of nose and mouth) appear in a darker tone of red, re-

vealing that the lighting variation imposed over the mannequin caused a greater luminance response over those surface normals with a smaller slant angle (forehead, cheeks, and continuous areas of the surface).

As far as the singular vectors are concerned, the element of the first left singular vector at 10° azimuth angle is depicted in (b) as \mathbf{u}_1 along with the weighted image $\sigma_1 \mathbf{u}_1$ (the singular vector scaled by its corresponding singular value). Similarly, the pair $(\mathbf{v}_1^T, \sigma_1 \mathbf{v}_1^T)$ is shown in (c), where images at 10° , 30° , and 60° zenith angles are depicted from top to bottom. The scaling of both left and right singular vectors by its corresponding singular value is shown in (d) as $\mathbf{u}_1 \sigma_1 \mathbf{v}_1^T$. The visual analysis of the pairs $(\mathbf{u}_1, \sigma_1 \mathbf{u}_1)$ and $(\mathbf{v}_1^T, \sigma_1 \mathbf{v}_1^T)$ reveals that the singular vectors encode information about which surface normals are illuminated or shadowed at a particular azimuth and zenith angles of the light source. Additionally, the singular vectors appear to include color information (red) around the shadowed areas of the mannequin while remaining in a rather neutral tone (gray) around the illuminated areas. The former observation indicates that, on one hand, the singular vectors of the luminance matrix are responsible of representing intrinsic features of the object such as color, specularity, and local shape of the sampled material. Singular values, on the other hand, determine aspects related with the light source and the effect it causes over the object's material. The color of the light source, the nature of the ambient light, the color of the surrounding objects, and inter-reflections might be mentioned amongst some of the features explained by singular values.

Results on fixing the zenith angle of the light source while varying the azimuth angle are shown in the bottom panel of Fig. 4. The first singular value is shown as a pixel map in (e), the pair $(\mathbf{v}_1^T, \sigma_1 \mathbf{v}_1^T)$ is shown in (f) for a fixed zenith angle of 30° , the pair $(\mathbf{u}_1, \sigma_1 \mathbf{u}_1)$ is depicted in (g) for variable azimuth angles of 120° , 190° , and 340° . The final weighting $\mathbf{u}_1 \sigma_1 \mathbf{v}_1^T$ is shown in (h). Similar observations as those made above can be derived from the visual analysis of the panel: (1) both the color and local shape of the object seem to be characterized by \mathbf{v}_1^T as the shadowed areas of the mannequin intensify the redness of the nail varnish

and (2) the singular values have a major influence around the illuminated areas, providing the color of the illuminated surface patch in accordance with the nature of the light source.

To conclude the analysis of Fig. 4, it is worth commenting that SVD appears to decompose a luminance matrix into orthogonal basis functions (singular vectors) spanning the reflectance properties inherent to the object's material, such as color, shininess, texture, and local surface orientation. These orthogonal basis functions are related through weighting factors (singular values) representing the characteristics inherent to the particular lighting with which the luminance matrices were created, such as color, intensity, and influence of the ambient and other lighting present at the moment of sampling. In the next section, it will be shown how this luminance matrix decomposition can be used in order to transfer reflectance between surface normals of different objects.

3 Correspondence and Transference

Each surface point is related to a reflectance history that can be transferred to another point with the same or very similar surface orientation. Thus, knowing the correspondences of the surface normals between the source and target objects is a key step. The correspondence process determines which surface normal locations are more similar between the source and target objects. This process can be achieved by exhaustively searching the normal maps of both objects. A normal map contains the surface orientation at each coordinate axis. However, this search results impractical and computational time consuming. We take advantage of the nature of Markov Random Fields (MRF) [5] and propose a method that statistically learns the surface normal correspondences between the source and target objects. We explain this method in what follows.

3.1 MRF Model for Surface Normal Correspondence

Our MRF model determines the correspondence distribution w_t in a searching area A_p of radius r . In this region, and by using a distance metric

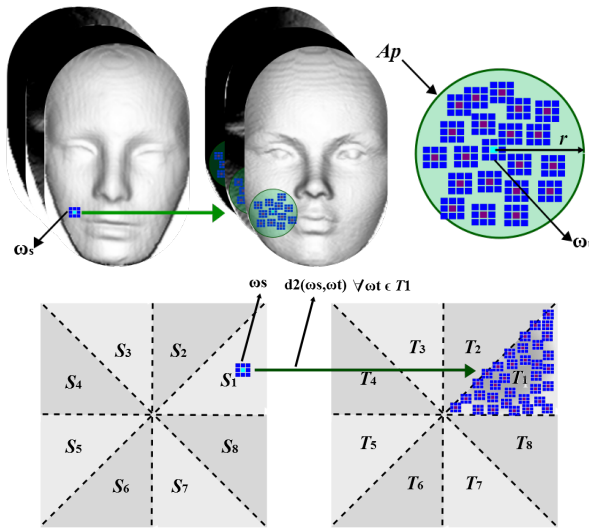


Fig. 5. Surface normal correspondence. A window of correspondence w_t is searched in the map of normals \mathbf{N}_T of the transfer object for each window w_s of the source object (face images). The region of search A_p is depicted by the circle (middle image) and is used to compute the probability distribution of the most probable window w_t that is similar to w_s . This type of correspondence works well for source and target objects that are geometrically similar. The squares at the right show the distribution of correspondence of normals based on its grouping. For each distribution w_s of sector S_i , the distribution of correspondence w_t is determined through the evaluation of a weighted distance d_2 of each distribution in T_i with w_s .

$d(w_t, w_s)$, we select the most similar neighborhood to w_s as the correspondence neighborhood w_t , i.e.,

$$w_t = \{w_t \in A_p : d(w_t, w_s) = 0\}; A_p \in \mathbf{N}_T. \quad (3)$$

where $w_s \in \mathbf{N}_S$ and $w_t \in \mathbf{N}_T$ are square windows centered in \mathbf{n}_s and \mathbf{n}_t , respectively. Subindexes s and t are used to refer to the source and target objects. Matrices \mathbf{N} represent the whole set of surface normals for each object, while \mathbf{n} are used to depict a single surface normal vector.

In an MRF model, the probability distribution of each normal \mathbf{n} , given the normals of its neighbors, is considered independent of the rest of locations in the normal map. The correspondence based on search regions requires *a priori* knowledge of some

geometric relation between the source and target objects. The reason for this is that the initial normal \mathbf{n}_t must be located in a geometrically similar region to the source object where \mathbf{n}_s is located. This correspondence works well for objects that are similar in geometry, such as spherical objects. However, as we want to be able to transfer the reflectance properties to objects with different geometry, we propose to group together the normals according to their surface orientation. Each normal has three components that correspond to its orientation at each axis coordinates. The normal orientations in the x and y axis contain values in the open interval $(-1, 1)$ and the orientation in the z axis is delimited by the interval $(0, 1)$. The grouping of normals allows to divide the normal maps \mathbf{N}_s and \mathbf{N}_t into eight sectors S and eight sectors T , respectively, according to the magnitude and sign of its components. For each neighborhood w_s in a given sector S_i , the corresponding neighborhood in T_i is determined by using a similarity distance between w_s with each neighborhood w_t in sector T_i (see the right image in Fig. 5). A weighted Euclidean distance metric between each pair of neighborhood is used to evaluate w_s with each neighborhood w_t .

This distance is described as follows: given $w_s \in S_i$ and some neighborhood $w_t \in T_i$, which contain three *layers* of components corresponding to the orientation in axis x , y , and z , matrices \mathbf{W}_s and \mathbf{W}_t are build such that each row is formed by two values of each layer, respectively. The dimension of each row vector is $k = k_m k_n$, where k_m and k_n represent the number of rows and columns of each neighborhood. The construction of matrices \mathbf{W} allows us to estimate a distance metric between w_s and w_t through the sum of the Euclidean distances between the row vectors of \mathbf{W}_s and \mathbf{W}_t .

Given that each surface point in each neighborhood is mutually related to its neighboring surface points, the local surface structure of the surface must be kept. This is achieved by assigning more weight to the normals that are close to \mathbf{n}_s and \mathbf{n}_t , and less weight to those that are far. However, the Euclidean distance does not maintain the surface locality as the values of each normal are considered independent. Furthermore, it is prone to variations due to surface scale. To preserve locality, a

Gaussian kernel $G(\sigma)$ is taken into account, which is contained in a vector \mathbf{g} of same dimension of row vector of matrices \mathbf{W} and weights the Euclidean distance, i.e.,

$$d_2(w_s, w_t) = \sum_{i=1}^3 \left(\begin{array}{c} g_1(\mathbf{W}_{s(i,1)} - \mathbf{W}_{t(i,1)})^2 + \\ g_2(\mathbf{W}_{s(i,2)} - \mathbf{W}_{t(i,2)})^2 + \\ \dots + g_k(\mathbf{W}_{s(i,k)} - \mathbf{W}_{t(i,k)})^2 \end{array} \right)^{\frac{1}{2}} \quad (4)$$

Equation 4 allows the evaluation of distance at each distribution $w_s \in S_i$ with all distributions w_t of the corresponding sector T_i . Each evaluation is then registered to a vector \mathbf{d} , such that the window of correspondence w_t is the window related to the minimum value of vector \mathbf{d} , i.e.,

$$\mathbf{d} = d_2(w_t, w_s); \forall w_t \in T_i, \quad (5)$$

$$w_t = \{w_t \in T_i : d_2(w_t, w_s) = \min(\mathbf{d})\}. \quad (6)$$

The normals of correspondence \mathbf{n}_s and \mathbf{n}_t will simply be the central elements of the neighborhoods of correspondence w_s and w_t , which are obtained from Equation 6. This process of correspondence is repeated for all the normals contained in sectors S of the map of normals of the source object \mathbf{N}_s . This way, the correspondence of normals between the source and transfer object is obtained.

3.2 Reflectance Transference

In order to calculate the new luminance matrices to be transferred onto the pixels of the target object, the original matrix \mathbf{M} is altered in accordance with the surface normal difference between \mathbf{n}_s and \mathbf{n}_t , where subindexes s and t respectively stand for source and target objects. As an example of the error calculation, the first singular value of \mathbf{M} is weighted by the absolute difference of the z -component between source and target surface normals as

$$\sigma'_1 = \sigma_1 \left(1 + w (|n_{sz}| - |n_{tz}|)^2 \right); \quad (7)$$

where $w = \begin{cases} 1; & |n_{sz}| < |n_{tz}| \\ -1; & |n_{sz}| > |n_{tz}| \end{cases}$

where n_{sz} and n_{tz} are respectively the z -component of the source and target surface normals, σ_1 is the first singular value of the luminance matrix, and σ'_1 is the first singular value of the new luminance matrix. The new second singular value, σ'_2 , is modified in accordance with the difference between the x -components of the surface normals, $(|n_{sx}| - |n_{tx}|)^2$. This similarly holds for the modified third singular value, σ'_3 , which is weighted using the y -component difference $(|n_{sy}| - |n_{ty}|)^2$. The new luminance matrices to be assigned to the pixels of the target object have now the form

$$\mathbf{M}' = \mathbf{U}\Sigma'\mathbf{V}^T;$$

$$\text{with } \Sigma' = \begin{pmatrix} \sigma'_1 & 0 & 0 & \dots & 0 \\ 0 & \sigma'_2 & 0 & \dots & 0 \\ 0 & 0 & \sigma'_3 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 0 \end{pmatrix}_{r \times r} \quad (8)$$

where the main diagonal of Σ' starts with the three modified singular values calculated using Eq. 7, containing zeros in the rest of the diagonal.

The fact that the first three singular values of the luminance matrices are individually modified using the z , x , and y components of the surface normals relies on two considerations. The first of these is concerned with the variability of the luminance matrix mostly retained by only a few basis vectors. The characterization of reflectance through low dimensional spaces has been widely studied [1, 9, 3], leading to the conclusion that reflectance of convex objects can be linearly approximated by a small number of principal components, in particular, the three first components suffice to characterize at least 90% of the appearance of the object under the imposed illumination changes. The second consideration is related with the influence of the lighting variability on the built basis functions [21], i.e., the singular vectors of a database considering photometric changes will tend to characterize the three main variabilities as orthogonal basis functions along the z , x , and y directions.

The effect of using Eq. 7 is explored in Fig. 6. Here, the surface normals of a Lambertian sphere




Fig. 6. Altering singular values of the luminance matrix. The figure illustrates the effect of transferring an altered luminance matrix in accordance with Eq. 8. In (a) a Lambertian sphere used as a source object is shown. The normals of the sphere are randomly rotated within $\pm 20^\circ$ in (b), i.e., becoming the target surface. Reflectance transference considering the difference error between surface normals is shown in (c). A further level of noise added ($\pm 45^\circ$) is depicted in (d) and its corresponding transference in (e). Note how not considering the surface normal error would have resulted not only in reflectance, but in shape transference as well




Fig. 7. Transferring reflectance onto synthetic spheres. Styrofoam mannequins covered with different substances such as a matte paint, sawdust, red nail varnish, and orange crayon are shown row-wise. Rendered views of a synthetic sphere after reflectance transference appear next to each mannequin image for illuminations with diffuse (left panel) and halogen (right panel) light sources. Note how appearance changes considerably when different light sources are applied on the same objects

are used as those of the source object. A rendered view of the Lambertian sphere is shown in (a). In order to create the target surface normal field, each source surface normal is randomly rotated within a $\pm 20^\circ$ interval, as shown in (b). Let us assume the pixel position (x, y) in the Lambertian sphere corresponds to the pixel position (x, y) in the disturbed sphere. If the correspondence error were not considered for altering the singular values

of the luminance matrix of the source object, not only the reflectance but also the shape would be transferred, i.e., the rendered views of the target object would be exactly the same rendered views of the source object. The result of modifying the luminance matrix is shown in (d), where it is clear the transference process takes into account the shape of the target object. A further level of added disturbance ($\pm 45^\circ$) is depicted in (d) and its




Fig. 8. The sphere as an intermediate reflectance object. Results of transferring reflectance from different materials onto the surface normals of a face from the Max Planck database are shown along the columns of the figure for two cases: (a) direct reflectance transference, (b) indirect reflectance transference using an intermediate sphere

corresponding transference in (e). Note how the quadratic term of Eq. 7 tends to punish greater correspondence errors between surface normals. This supports the idea of using SVD for luminance matrix decomposition as opposed to using the original values of the luminance matrix directly, since altering reflectance at the pixel level would be difficult otherwise.

3.3 A Sphere as a Transference Object

Fig. 7 presents results for transferring reflectance of different materials onto the surface normals of a synthetic sphere. The figure contains two panels. The left panel shows results for the objects illuminated by a diffuse light while the right panel shows results for illumination with a halogen light. Note how each light is differently reflected by the same object, i.e., the color of the mannequin changes when using different light sources. This makes, for example, the matte painting look ochre under diffuse lighting and reflect an orange tonality under halogen lighting. As previously discussed in Section 2, it was observed that reflectance features caused by the nature of the light source are transferred onto the surface normals of the spheres through the singular values of its luminance matrices. The weighting process carried out in Eq. 7 assigns a transference confidence depending on the discrepancy between corresponded surface normals. It is also noticeable how the specular

behavior mainly exposed by the crayon and varnish materials is conserved after luminance matrix transference. Similar observations hold for the diffuse reflectance of the matte paint and the scattered texture of the sawdust.

The importance of using a sphere as an intermediate reflectance object is illustrated in Fig. 8, where transference results on a face of the Max Planck database [2] are shown. The top row of the column presents rendered views of the face after direct transference, i.e., performing surface normal correspondence directly between the surface normals of the different mannequins and the surface normals of the face model. The bottom row of the figure depicts results after using a sphere as an intermediate reflectance object, i.e., luminance matrices are first transferred from the mannequin models to a synthetic sphere (as in Fig. 7), then a new transference process is performed using the sphere as the source object. The visual comparison of the rows of the figure demonstrates a benefit when the sphere is used to transfer reflectance between the two gradient fields. Specifically, regions where the surface normals are subject to change in direction, such as edges around the face, appear more defined in the results shown in the bottom row, diminishing abrupt transitions between neighboring pixels. Interestingly, the main differences between the images corresponding to the crayon material are located around the eyes, eyebrows, and mouth areas, where the crayon

mannequin was painted on purpose with a black crayon. For direct transference, it is evident that the local correspondence between the facial areas of the mannequin and the face model are responsible for the black crayon being assigned onto the pixels of the target object. This effect is nonetheless suppressed after using a sphere as an intermediate source object.

4 Experimental Results

This section presents experimental results for the proposed reflectance transference methodology. A sphere sized 150×150 pixels was applied in all the experiments as an intermediate reflectance object. Also, for all experiments, three target objects were used from the Stanford database [24]: rabbit (150×150 pixels), dragon (150×200 pixels), and Buddah (200×80 pixels). The results are organized in panels, each of them presenting in its first row the source object (mannequin) as illuminated by different light source directions. The remaining rows present reflectance transference for the target objects. Both scenarios using the diffuse and halogen lighting are depicted in the figures.

Transference results for the matte paint and the sawdust mannequins are illustrated in Fig. 9. The first feature to note from the figure is that the reflected color seems to be accurately transferred onto the Stanford surfaces. The subtle difference in the orange tone provoked by the halogen light between the matte paint and the sawdust mannequins is also noticeable along the rendered views of the target objects. In a similar way, this phenomenon is observable for the diffuse lighting case, where the difference between the ochre and the golden tones for the matte paint and the sawdust is also present after reflectance transference. Another feature to note from the figure is the scattered texture of the sawdust, which is passed onto the surface of the target objects.

Fig. 10 presents the specular reflectance cases, showing results for the nail varnish and crayon mannequins. Again, the tones of orange and red color are evidently transferred over the target objects. The strong specularities of the nail varnish appear as well strengthened along the novel rendered views, particularly for the halogen lighting.

Also, the wax-like consistency of the crayon is noticeable, especially for the diffuse light renderings.

It is worth commenting on the presence of ambient light in the source database. For the diffuse lamp, the data acquisition was performed in a closed dark room, therefore, the only illumination affecting the mannequins was the diffuse light itself and, as a consequence, pronounced shadowed regions appear over the surface of the illuminated mannequins. On the contrary, the halogen lamp was subject to the influence of ambient light, as experiments were performed in our lab, where artificial and natural light coming from open windows intervened in the sampling process. In this case, transference results obtained using the halogen lamp reveal more illuminated regions around the boundaries of the mannequins, as opposed to the shadowed regions exhibited by the examples corresponding to the diffuse lamp setting. This observation is corroborated in the transferred reflectance of the Stanford objects, i.e., the ears of the bunny appear shadowed for the diffuse lighting, while this effect is lessened for the halogen lighting.

Figures 9 and 10 present individual images from both the source and transferred databases, allowing a qualitative examination of the transference results obtained from the method proposed in this paper. Results at the global level, nonetheless, are necessary to assess similarity between databases in a numerical way. To this end, we chose a methodology similar to the parametric eigenspace of [19], where each image is projected onto a considerably lower dimensional space in order to analyze the behavior of the image ensemble in the feature space spanned by the database. The experiment is described as follows: for each image, a long column vector was created by stacking the image columns in order to create a big matrix \mathbf{A} , whose SVD was later applied for calculating singular values and singular vectors as

$$\mathbf{A}_{p \times n} = \mathbf{U}_{p \times 3} \mathbf{\Sigma}_{3 \times 3} \mathbf{V}_{3 \times n}^T. \quad (9)$$

where p is the number of image pixels and n is the number of images in the database. The columns of \mathbf{V}^T contain the projection of the n database images in R^3 . Each database was decomposed separately for the R, G, and B channels, and results of this decomposition are shown in Fig. 11.

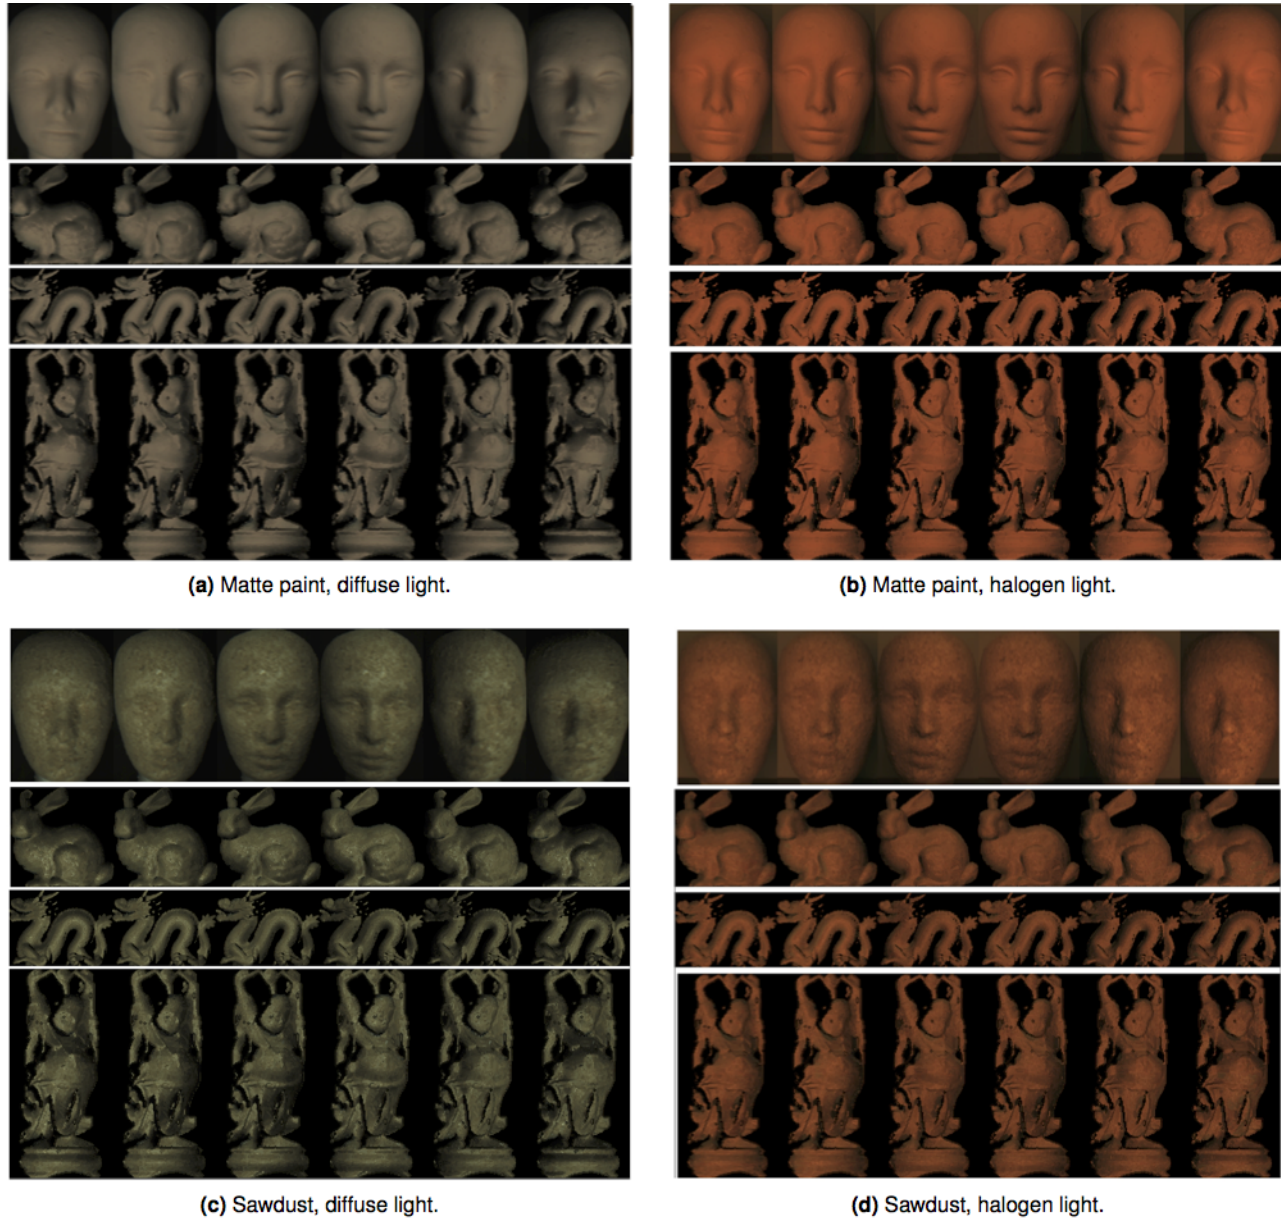


Fig. 9. Reflectance transference for matte and scattered surfaces. The figure illustrates how the reflectance of a styrofoam mannequin is transferred onto the 3D shape of the Stanford database rabbit, dragon, and Buddha. The mannequin was illuminated using a diffuse and a halogen light. Corresponding rendered views after reflectance transference are shown in (a) and (b) for the matte paint, while the results for transferring sawdust reflectance are shown in (c) and (d)

The retained variability for the databases is shown in (a), where four diagrams depict the cumulative singular values for the different sampled materials:

matte paint, red varnish, orange crayon, and sawdust. Each of the four diagrams is related to one single material, presenting cumulative plots for the

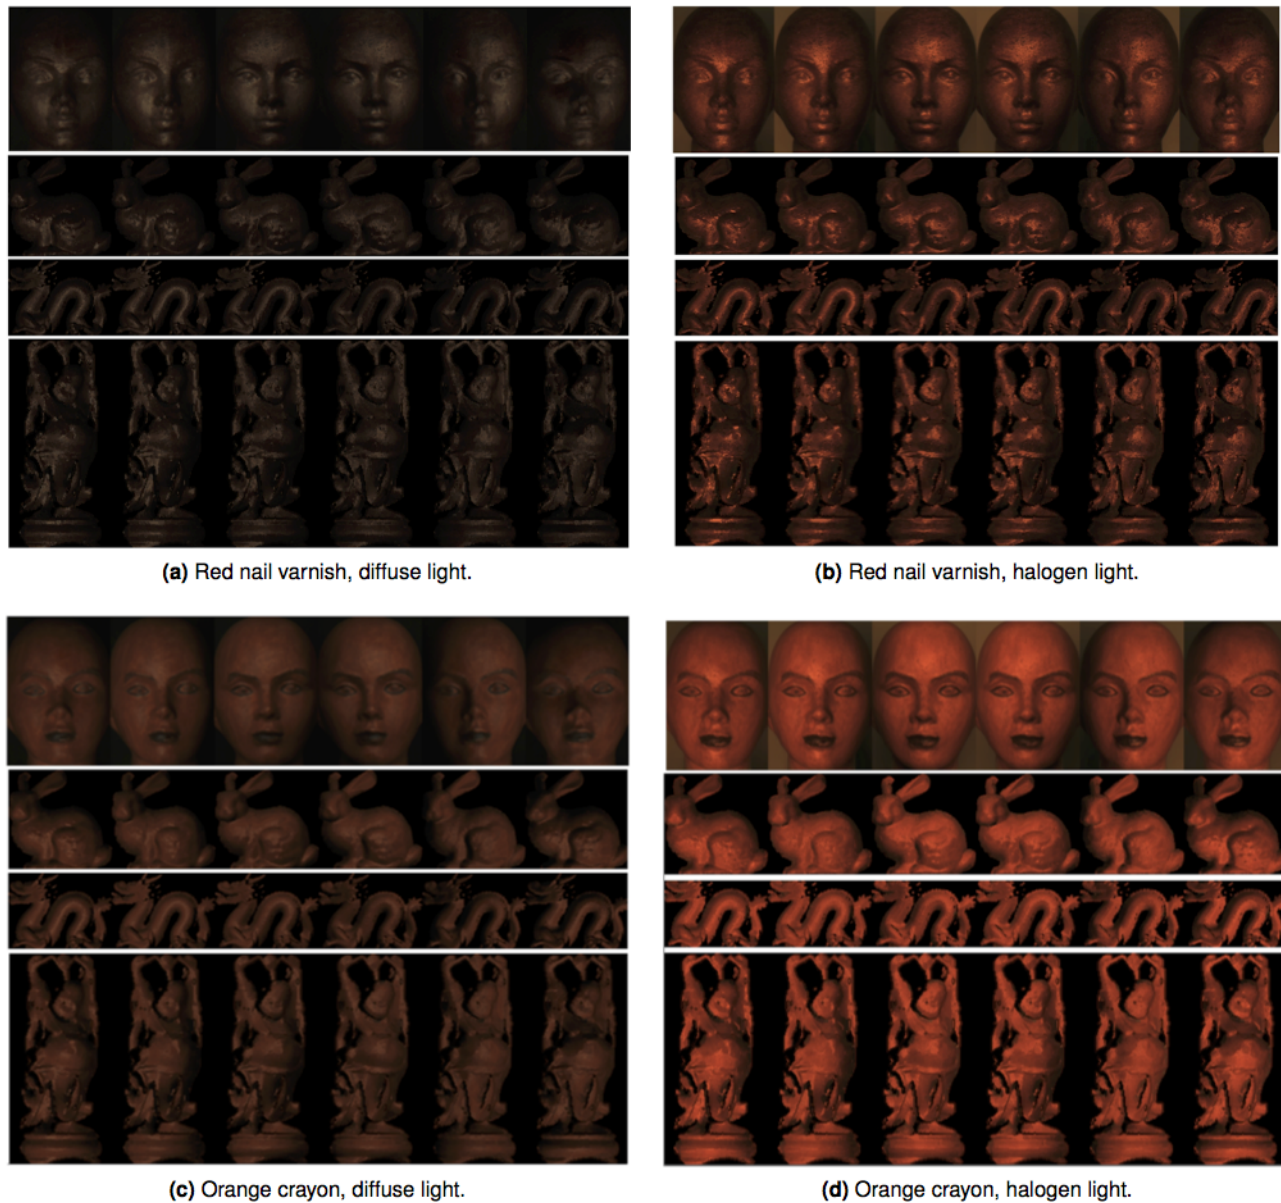


Fig. 10. Reflectance transference for specular surfaces. The figure illustrates how the reflectance of a styrofoam mannequin is transferred to the 3D shape of the Stanford database rabbit, dragon, and Buddha. The mannequin was illuminated using a diffuse and a halogen light. The corresponding rendered views after reflectance transference are shown in (a) and (b) for red nail varnish, while the results for transferring orange crayon are shown in (c) and (d)

four objects (mannequin, bunny, Buddha, dragon), three color channels (R, G, B) and two illumination sources (diffuse, halogen) for a total of $4 \times 3 \times 2 = 24$ line plots for each material. The singular values

of the databases acquired using halogen light are shown with blue lines, while those acquired under diffuse illumination are presented in red. The intention of the plots is to show how the databases can

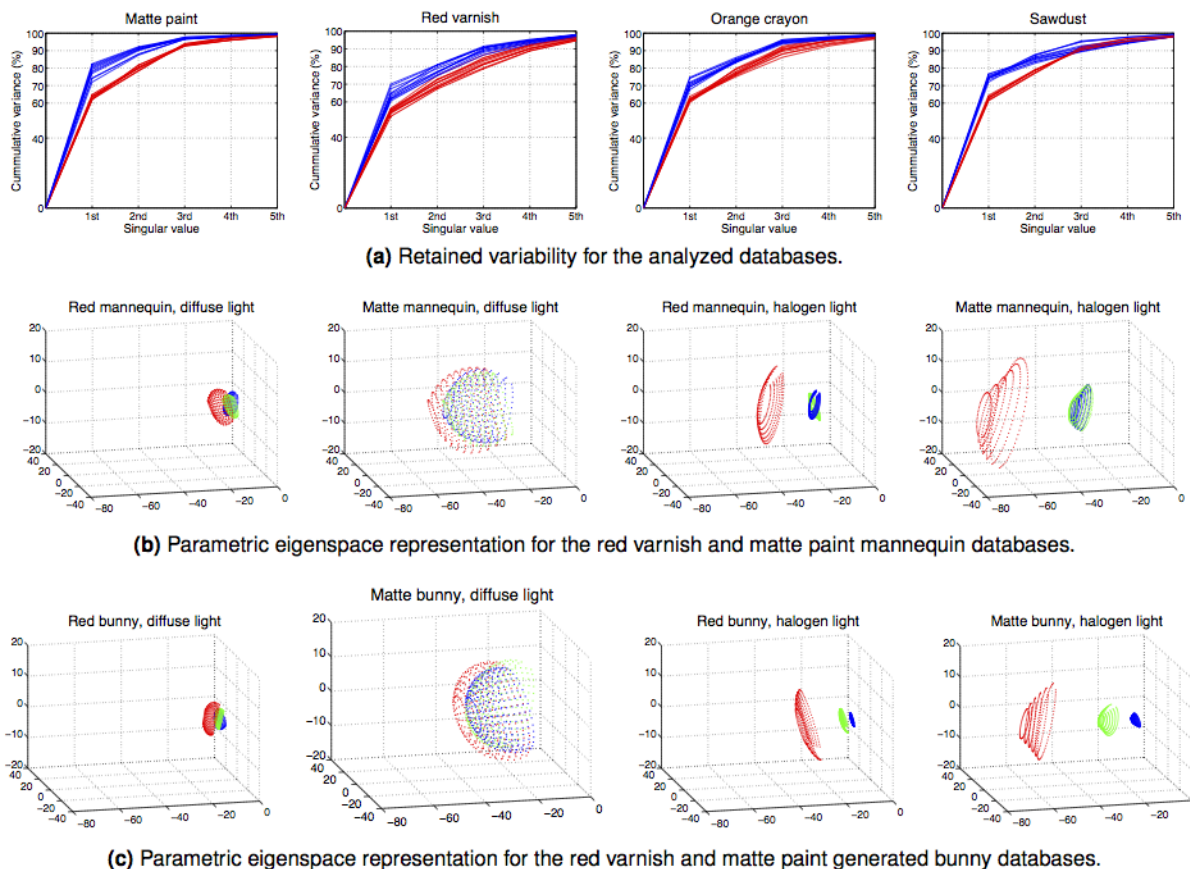


Fig. 11. Reflectance analysis in the parametric eigenspace. In (a) the retained variability of the original mannequin databases and synthetically generated databases is shown for the halogen (blue) and diffuse (red) lines. The purpose of these diagrams is to show how the databases cluster depending on the light source used in the original experiment, indicating that this information has been passed onto the target objects after transference. The results of projecting each image of the database onto the parametric eigenspace are depicted in (b) for the red varnish and matte mannequin databases (natural reflectance) and (c) for the reflectance transferred bunny databases. Each image is shown in colored dots, where the color refers to database decomposition for the red, green, and blue color channels

be clustered mainly in accordance with the type of illumination used to create the source database, regardless of the shape of the object or whether the database was generated naturally or by reflectance transference.

The remaining rows of Fig. 11 show the spatial arrangement of the eigenvectors of the databases, i.e., each image in the database is projected onto the three dimensional eigenspace using $\mathbf{v}_i^T \Sigma$, $i = \{1, 2, \dots, n\}$, where the row vector $\mathbf{v}_i^T_{1 \times 3}$ is a column of the right singular matrix \mathbf{V} . The results of this projection are depicted in (b) for the red

varnish and matte mannequin databases (natural reflectance) and the reflectance transferred bunny databases appear in (c). In all the diagrams, the projected images are shown in colored dots, where the color refers to database decomposition for the red, green, and blue color channels. The visual analysis of the plots reveals a resemblance between the natural and the transferred reflectance behaviors, suggesting that both original and synthetically generated databases span the similar illumination subspaces, corroborating the success of the reflectance transference process.

5 Conclusions

A method for characterization and transference of reflectance has been proposed in this paper. Ideas borrowed by extended photometric sampling, which decomposes luminance matrices into singular vectors and singular values are used to represent reflectance variations of a source object acquired by a camera and a moving light source. A method based on Markov Random Fields is also introduced to solve the correspondence problem between the surface normals of a source object and the surface normals of a sphere, which is later used to transfer reflectance of the source object onto a target object. Experiments demonstrate that our method successfully transfers reflectance behavior for a variety of materials and under different illumination settings. Given that the statistical decomposition of reflectance allows a further analysis, the parametrization of the luminance matrices, i.e., through the manipulation of the singular values is considered as future work for generating synthetic reflectance.

References

1. **Basri, R. & Jacobs, D. (2003).** Lambertian reflectance and linear subspaces. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 25, No. 6, pp. 383–390.
2. **Blanz, V. & Vetter, T. (1999).** A morphable model for the synthesis of 3d faces. *Proc. SIGGRAPH*, volume 1, pp. 187–194.
3. **Chen, H., Belhumeur, P., & Jacobs, D. (2000).** In search of illumination invariants. *Proc. IEEE International Conference in Computer Vision and Pattern Recognition*, pp. 1–8.
4. **Cook, L. R. & Torrance, K. E. (1982).** A reflectance model for computer graphics. *ACM Trans. Graph.*, Vol. 1, pp. 7–24.
5. **Cross, G. R. & Jain, A. K. (1983).** Markov random field texture models. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, Vol. 5, pp. 25–39.
6. **Curet (1999).** *Columbia-utrech Reflectance and Texture Database*. [Http://www.cs.columbia.edu/CAVE/curet/](http://www.cs.columbia.edu/CAVE/curet/).
7. **Dana, K. J. (2001).** Brdf/btf measurement device. *Proc. IEEE International Conference in Computer Vision*, volume 2, pp. 460–466.
8. **Dana, K. J., Ginneken, B. V., Nayar, S. K., & Koenderink, J. J. (1999).** Reflectance and texture of real-world surfaces. *ACM Trans. Graph.*, Vol. 18, pp. 1–34.
9. **Epstein, R., Hallinan, P. W., & Yuille, A. L. (1995).** 5 ± 2 eigenimages suffice: an empirical investigation of low-dimensional lighting models. *Proc. Workshop on Physics-based Modelling in Computer Vision*, pp. 108–116.
10. **F. E. Nicodemus, I. W. G., J. C. Richmond & Limperis, T. (1977).** Geometrical considerations and nomenclature for reflectance. *NBS Monograph*.
11. **Ghosh, A., Heidrich, W., Achutha, S., & O'Toole, M. (2010).** A basis illumination approach to BRDF measurement. *International Journal of Computer Vision*, Vol. 90, pp. 183–197.
12. **He, X. D., Torrance, K. E., Sillion, F. X., & Greenberg, D. P. (1991).** A comprehensive physical model for light reflection. *Proc. SIGGRAPH Computer Graphics*, volume 25, pp. 175–186.
13. **Hernández-Rodríguez, F. & Castelán, M. (2012).** A method for improving consistency in photometric databases. *Proc. British Machine Vision Conference*, pp. 1–10.
14. **Hertzmann, A. & Seitz, S. M. (2005).** Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 27, No. 8, pp. 1254–1264.
15. **Liu, R. & Han, J. (2010).** Recovering surface normal of specular object by hough transform method. *IET Computer Vision*, Vol. 4, No. 2, pp. 129–137.
16. **Marschner, S. R., Westin, S. H., Lafortune, E. P. F., & Torrance, K. E. (2000).** Image-based bidirectional reflectance distribution function measurement. *Appl. Opt.*, Vol. 39, No. 16, pp. 2592–2600.
17. **Matusik, W., Pfister, H., Brand, M., & McMillan, L. (2000).** A data-driven reflectance model. *ACM Trans. Graph.*, Vol. 22, No. 3, pp. 759–769.
18. **Mertens, T., Kautz, J., Chen, J., Bekaert, P., & Durand, F. (2006).** Texture transfer using geometry correlation. *Proc. Eurographics Symposium on Rendering*, pp. 273–284.
19. **Murase, H. & Nayar, S. (1993).** Parametric eigenspace representation for visual learning and recognition. *Proc. SPIE Geometric Methods in Computer Vision II*, volume 2031, pp. 378–391.

20. **Phong, B. T. (1975)**. Illumination for computer generated pictures. *Communications of ACM*, Vol. 18, No. 6, pp. 311–317.
21. **Ramamoorthi, R. (2002)**. Analytic pca reconstruction for theoretical analysis of lighting variability in images of a lambertian object. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 24, No. 10, pp. 1322–1333.
22. **Saito, H., Omata, K., & Ozawa, S. (2003)**. Recovery of shape and surface reflectance of specular object from relative rotation of light source. *Image and Vision Computing*, Vol. 21, pp. 777–787.
23. **Sato, I., Okabe, T., Yu, Q., & Sato, Y. (2007)**. Shape reconstruction based on similarity in radiance changes under varying illumination. *IEEE International Conference in Computer Vision*, pp. 1–8.
24. **Stanford.Repository (2011)**. The Stanford 3D scanning repository. <http://graphics.stanford.edu/data/3Dscanrep/>.
25. **Ward., G. J. (1992)**. Measuring and modeling anisotropic reflection. *Proc. SIGGRAPH Comput. Graph.*, volume 24, pp. 265–272.

Mario Castelán obtained his Ph.D. in Computer Science from the University of York, U.K., in 2006. Currently, he is a full-time researcher at the Robotics and Advanced Manufacturing Research Group of CINVESTAV - Saltillo. His research interests are focused on 3D shape analysis and

statistical learning for data analysis, computer vision, and robotics applications.

Elier Cruz-Pérez obtained the M.Sc. in Robotics and Advanced Manufacturing in 2013 from CINVESTAV, Mexico. His research interests are in computer vision, machine learning, and advanced manufacturing.

Luz Abril Torres-Méndez received her Ph.D. degree from the McGill University, Canada, in 2005. Since 2006 she has been a Full Professor at the Robotics and Advanced Manufacturing Group of the Research Centre and Advanced Studies (CINVESTAV) Campus Saltillo, Mexico. She is the Founder and Head of the Vision for Robotics Lab in the same institution. Her primary research interests are computer vision and mobile robotics: in particular, 3D environment modeling, active perception, underwater image enhancement, super-resolution, color and surface normal synthesis on faces; human-robot interaction, transferring cognitive skills to robots, sensor fusion, vision-based navigation for wheeled and underwater robotics.

*Article received on 20/02/2014; accepted on 06/05/2015.
Corresponding author is Mario Castelán.*