Light calibration and quality assessment methods for Reflectance Transformation Imaging applied to artworks' analysis

A. Giachetti[1*], C. Daffara[1], C. Reghelin[1], E. Gobbetti[2], R. Pintus[2]
1) Dept. of Computer Science, University of Verona, Strada Le Grazie, 15, 37134 Verona (IT)
2) Visual Computing Group, CRS4, Italy

ABSTRACT

In this paper we analyze some problems related to the acquisition of multiple illumination images for Polynomial Texture Maps (PTM) or generic Reflectance Transform Imaging (RTI). We show that intensity and directionality non-uniformity can be a relevant issue when acquiring manual sets of images with the standard highlight-based setup both using a flash lamp and a LED light. To maintain a cheap and flexible acquisition setup that can be used by non experienced users we propose to use a dynamic calibration and correction of the lights based on multiple intensity and direction estimation around the imaged object during the acquisition.

Preliminary tests on the results obtained have been performed by acquiring a specifically designed 3D printed pattern to see the accuracy of the acquisition obtained both for spatial discrimination of small structures and normal estimation, and on samples of different types of paper to evaluate material discrimination.

We plan to derive from our analysis and from the tools developed and under development novel procedures and guidelines that can be used to turn the cheap and common RTI acquisition setup from a simple way to enrich object visualization into a powerful method for extracting quantitative characterization of surface geometry or of reflective properties of different materials. This results could have relevant applications in the Cultural Heritage domain, in order to recognize different materials used in paintings or ageing status of artifacts.

Keywords: Reflectance Transformation Imaging, Polynomial Texture Maps, calibration, surface reconstruction, artworks materials.

1. INTRODUCTION

Reflectance Transformation Imaging (RTI) is becoming a popular tool for the acquisition of paintings and other kinds of artworks, due to the possibility of performing the acquisition with a simple setup based on a camera, a single light source and a reflective sphere for the estimation of the illumination direction [1,2]. This kind of setup is quite low cost, flexible, and portable, being therefore particularly suitable for the acquisition of large surfaces of artworks e.g. mural paintings and bas-reliefs on site.

However, the results obtained with this simple setup may be not suitable for the quantitative analyses derived from the RTI reconstruction and proposed in the literature - as normal estimation, 3D reconstruction, material classification [3,4] due to lack of uniformity in the illumination intensity and direction.

This fact is a relevant limit for the possible exploitation of the acquired data because the accuracy in the recovery of details and the possibility of having a good characterization of the reflective properties of materials are extremely important in Cultural Heritage applications, as they could give relevant hints for the historical analysis of the artworks and to understand and recognize different painting techniques. Errors in the assumptions regarding the illumination pattern may also create artefacts in the shape detail reconstruction and may affect the correct interpretation of the data. If we want to derive from coefficients of polynomial (PTM [5]) or Spherical Harmonics-based (HSH [2]) fitting multi-light acquisition, intrinsic properties of materials or geometric information (normals that can be used for surface recovery), we should assume to have at each pixel location a correct estimation of light intensity and direction (and also to know if there are shadows or specular highlights [7,8]).

*andrea.giachetti@univr.it phone +39 045 8027998 fax +39 045 8027068 ; www.univr.it

Figure 1: Left: Empty acquisition setup with white surface and four reflective spheres. Right: Setup with example objects (paper samples and 3D printed target

These assumptions are clearly not valid in common acquisitions performed on artworks and other artefacts in the common practice, for several reasons. Light projected by the lamps used at a relatively short distance is not uniform, both in intensity and direction, and its position and direction is not accurately controlled. We will show the amount of error measured in a simple setup in Section 3, after a short introduction of the RTI technique. In Section 4 we will show that using simple correction technique not requiring using special lights or preliminary calibration procedures it is possible to improve the quality of the PTM reconstructions obtained from data.

2. REFLECTANCE TRANSFORMATION IMAGING

Reflectance Transformation Imaging (RTI) is an imaging technique originally developed at HP labs [5] and based on acquisition of multiple images with fixed camera and moving lights of a target object approximately perpendicular to the camera axis. Assuming that lights intensity is constant and images are acquired with a sufficient number of well sampled light directions, it is possible to fit a parametric model over the image stack to encode per pixel color and luminance as a function of arbitrary light direction. In the classical Polynomial Texture Maps (PTM) model, the parametric function is a second order polynomial of the x,y light direction projection components, assuming constant chromaticity, as in equations (1)

$$R(x,y)=R_0(x,y)*L(x,y,l_x,l_y)$$
$$G(x,y)=G_0(x,y)*L(x,y,l_x,l_y)$$
$$B(x,y)=B_0(x,y)*L(x,y,l_x,l_y)$$

$$L(x,y,l_x,l_y)=a_0 l_x^2+a_1 l_y^2+a_2 l_x l_y+a_3 l_x+a_4 l_y+a_5 \tag{1}$$

This allow the use of specialized interface as RTI viewer [9] to perform an off line analysis of the object of interest with interactive relighting to easily inspect all its detail, possibly using special enhancement functions using multiple illumination information to enhance rendered images. This is clearly a particularly interesting technique for artworks analysis, as it can be seen as a way to extend and use at home the raking light visualization typically applied by conservators and restorers.

More information can then be obtained by the RTI files, as the parametric representation can be used to estimate normals and reconstruct 3D shape (photometric stereo techniques). However, the 3D reconstruction based on multiple illumination requires accurate control of light intensity and direction (an analysis of the error propagated by inaccuracy is presented in [6]). The reconstructed function L can be considered also a slicing of the material BRDF [3], allowing ideally material discrimination from PTM/RTI reconstructions. But also in this case, the local characterization makes sense only if a limited error in the constant illumination assumption is actually obtained in the acquisition phase.

Limitation of the error is, however quite hard to be obtained even in complex dome setups. Most of the RTI acquisitions applied in the Cultural Heritage domain are then performed using the classical setup based on manual light positioning and on the use of a reflective sphere to recover the light direction vector.

In order to understand if it is possible to obtain a quantitative analysis of RTI reconstruction with this kind of setup, we performed some experiments measuring light inhomogeneity and proposing simple methods to correct the acquisition errors.

3. ACQUISITION SETUP AND EXPERIMENTAL DESIGN

In order to evaluate illumination direction and intensity we created the following setup: a flat surface covered with a uniform approximately Lambertian paper sheet, with four reflective black spheres used for the classical highlight based light direction estimation (Figure 1) at the corners. We have built a simple semicircular support in order to easily move the light over an approximately spherical surface of about 1 m of radius around the center of the sheet. However, the light is manually placed in the desired position, so that an accurate position and direction cannot be determined as in the classical highlight-based RTI setups.

In our experiments we used two different light sources: a flash light synchronized with the camera and a 10 W LED spot..

The camera was fixed at 1m of distance from the sheet with optical axis perpendicular to the flat surface. We also calibrated the camera setup finding a negligible distortion error with the fixed parameters.

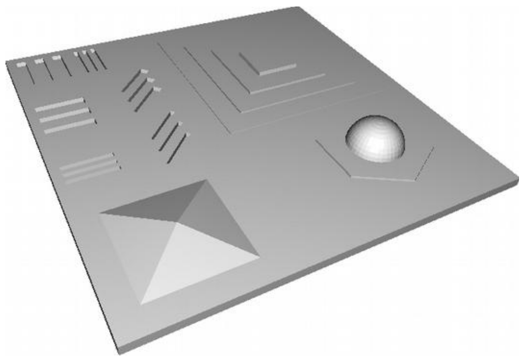

Figure 2: Design of the 3D printed target used to characterize geometrical reconstruction quality

1. This setup was designed with two goals: to characterize the light intensity and direction non-uniformity by estimating image intensity variation along the white surface and inconsistencies in directions estimated at the four sphere locations
2. to create a setup with an associated procedure to acquire “corrected” RTI of objects, on the basis of the non-homogeneities evaluated at the margins of the object support area. Our idea is to acquire objects in the central part of the image, surrounded by a planar white Lambertian frame to estimate illumination and four spheres at the frame corners for independent estimation of direction.

To evaluate the quality of the reconstructions obtained we used two different targets: a 3D printed object, and a set of different materials over a planar support.

The 3D printed object (Figure 2) was designed to evaluate normal inaccuracies but also to evaluate sharpness in discontinuity discrimination. We put over a plane a small sphere (R=6rmm) a pyramid (side=2cm), differently oriented slots (min thickness 1mm) and different steps with heights starting from 0.2 mm that is the accuracy of the printing procedure used for the physical realization of the current shape target. The material chosen for the printing was an opaque plastic with low specular effects. The printing procedure created small but visible staircase artifacts on the pyramid surface.

The different materials were simply pieces of papers with different reflectance and roughness surface micro-structure attached on the support.

4. TEST ACQUISITIONS, ERROR ANALYSIS AND IMAGE CORRECTIONS

In the first test we moved the lights on an arc on the left of the camera keeping it at approximately 1 meter of distance from the center of the plane and directed towards it.

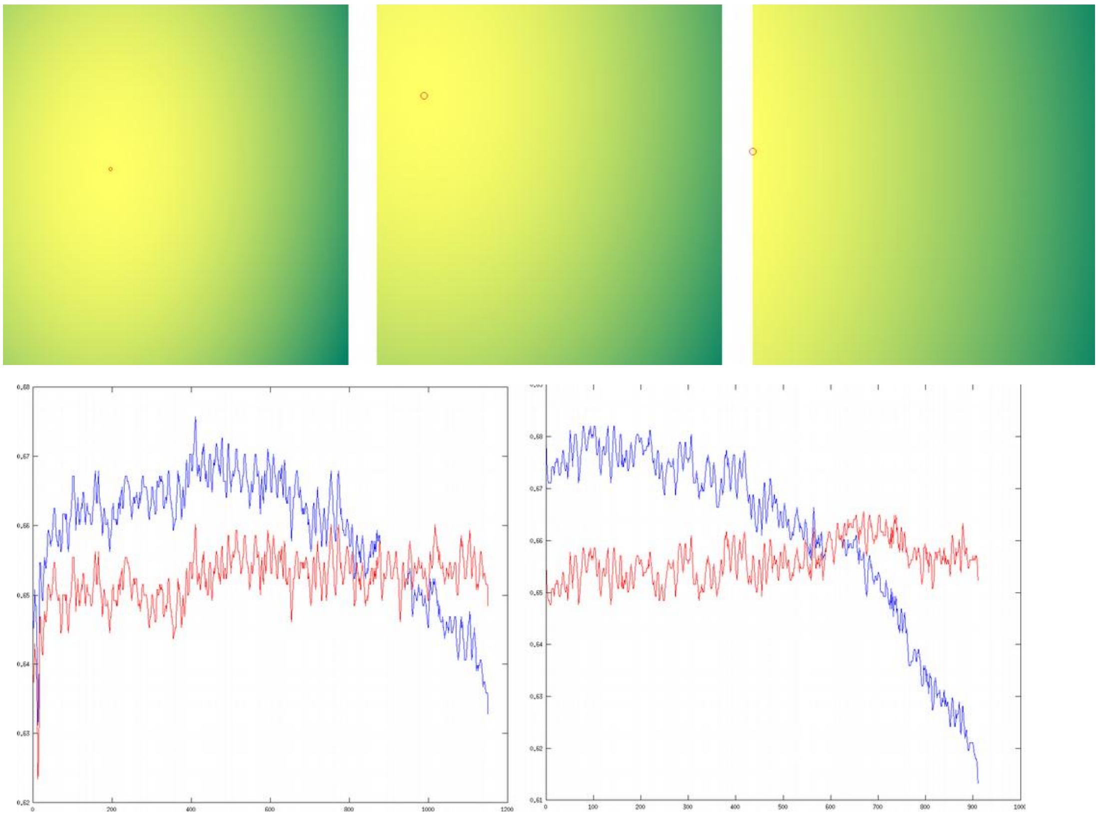

Figure 3: Profiles of the light intensity measured on the x and y axes (blue lines). The red lines represent the profiles after quadratic correction derived from a polynomial fit on values measured at plane borders.

Looking at the variability of the intensity over the surface, we found some interesting results: the light intensity measured by the camera sensor on the white surface was not constant, and the actual light pointing is not easily controlled by operators. These facts are quite evident looking at Figure 3. Images on the first row represent a color-mapped quadratic interpolation of the intensity measured on the white plane with a dot showing the location of the maximum of the function. Differences are enhanced by the colormap, however it is possible to see that, assuming that the maximum should correspond to the center of the target, the light pointing is not accurate, as it may be expected as the light is manually adjusted and that the light is not homogeneous. Here the results are obtained using the flash light and different elevation values. Similar results were obtained with the LED light spot, with an evident inaccuracy in light pointing and a varying light intensity.

In the second row od Figure 3 we represented with blue lines the intensity profiles on the white surface along x and y axes in the reference centered in the mid point of the image for a fixed value of the elevation. The lines show clearly that there are both effects depending on light spot shape and attenuation with distance. Plotting the average intensity on the white surface as a function of angular elevation we found that it however, approximately followed the cosine law as expected (Figure 4 left).

It is clearly not viable for practical and low cost RTI acquisition to look for a perfect light direction alignment and the use of special lights, as it is not possible in general to move light sources too far from the object. So, our idea is to automatically calibrate/correct the images acquired with bad lights and pointing control, exploiting information captured by targets placed around the object of interest. The procedure we designed and implemented is the following: we sample the intensity along the planar Lambertian frame (the white rectangle border in our experiments). We take 12 samples locally averaging the intensity in 5x5 windows). A quadratic function of image coordinates is then fitted over samples. Applying a multiplicative correction in order to flatten the estimated profile while keeping the average intensity we obtain intensity corrected images without preliminary light characterization. Red plots in Figure 4 are the corrected

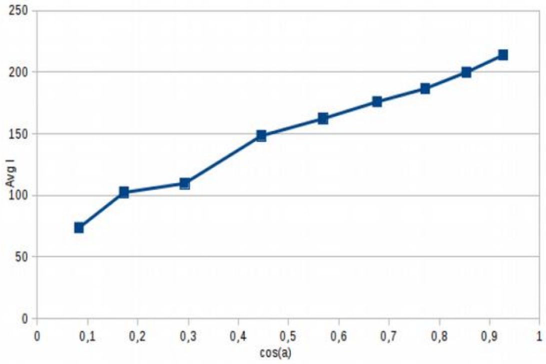

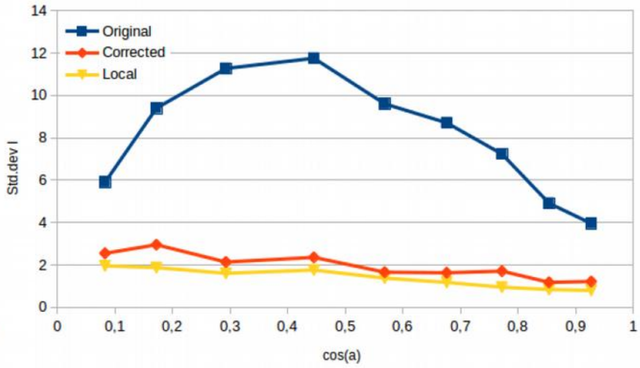

Figure 4: Left: average intensity on the white plane versus light elevation for the Flash light test. Right: white plane intensity variance without (blue line/squares) and with (red line/diamonds) correction as a function of light elevation (Flash light). Yellow line represents the local variance due to the paper texture.

image profiles corresponding to the original ones and demonstrate that the resulting illumination pattern is approximately constant.

The effect of the correction is also visible in Figure 4 (right), representing the intensity variance in the white plane vs angular elevation for the Flash light test. The original values (blue line/squares) show a non negligible error and a behavior that depends on the intensity source features and on direction accuracy. Corrected images, on the other hand have a variance that is reduced to a consistently smaller value no longer dependent on the elevation of the light source. Its value is quite close to the one estimated in a local 40x40 patch that we assume to be due to the paper texture and is represented by the yellow line with triangular markers. This means that the quadratic correction is effective in eliminating the intensity inhomogeneity on the entire image and as it was estimated on the border area, it can be applied as well for acquisitions with the objects of interest in the central area. Similar results were obtained with the LED light spot.

If we consider the inhomogeneity of light direction on the plane surface, we can look at the differences in the direction estimated on the four spheres. In this case we acquired images sampling uniformly the hemisphere as doing a classical RTI acquisition. We estimated the light direction on the four spheres with the functions available in the RTI builder package, using manual correction of the sphere contour. Spheres in the preprocessed images have an apparent diameter of 101 pixels. The fact that the assumption of constant light direction is violated can be seen by measuring the differences in the directions measured on the different spheres.

Figure 5 shows a plot with the difference between directions measured at diagonally opposite spheres, plotted against the corresponding angular elevation of the light source (blue squares). The error is quite high, much more than the expected quantization error that also changes with the elevation due the increased uncertainty in the highlight position (represented by red diamonds).

In PTM estimation we do not necessarily need to assume constant light direction on the plane as we actually fit independently the model on each pixel. We cannot estimate the direction everywhere, but we can interpolate from the multiple spheres. Yellow dots in Figure 6 actually show the difference between the value computed on a sphere and the value obtained from a linear interpolation of the three directions estimated on the other three spheres in the same location. This comparison is done as it measure the difference between the worst-case error in the direction estimation that we have when using the same direction measured on one sphere for all the pixels to obtain the PTM fit, and the worst-case error we would have if for each pixel we compute a local direction interpolated from multiple spheres. It is possible to see that the error is drastically reduced.

This means that, also for light direction, provided we can place multiple reflective spheres around the object, we can obtain a per-pixel corrected light direction estimate from the single estimates and recover a possibly more effective reconstruction of RTI parameters.

Note that for the directional correction this means, it is necessary to modify the currently used PTM/HSH fitting tools in order to support a locally adapted light direction. We are currently developing novel versions of the fitters, also taking into consideration the problem of removing or segmenting projected shadows/specular highlights.

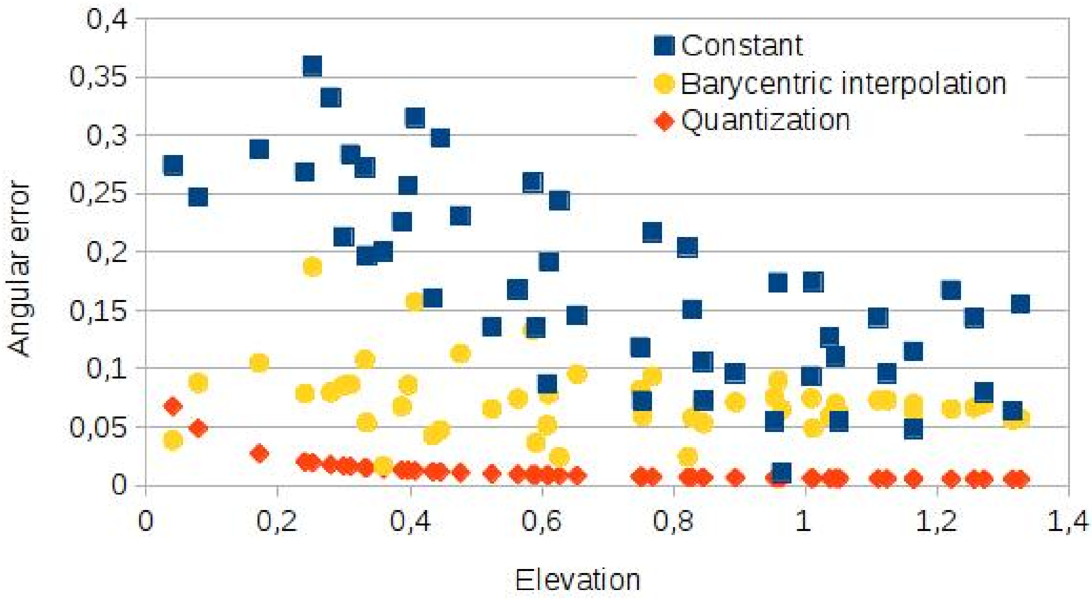

Figure 5: Differences between direction measurements at opposite spheres (top left/bottom right) for the Led light (Blue squares) as a function of the angular elevation. Different values for similar elevation depend on different azimuthal angles. Yellow dots are the differences between top left sphere values and barycentric interpolation derived from the other three spheres. Red diamonds represent the expected errors due to the pixel quantization.

5. LIGHT INTENSITY CORRECTION AND PTM

Intensity correction seems actually to improve both the quality of the 3D reconstruction and the material characterization. Figure 6 shows the normal reconstruction obtained from PTM coefficients computed on the images of the printed target illuminated with flash light without and with corrections. The normal image obtained from the corrected ones appears sharper and the normal plot over the sphere presents clearly reduced artifacts.

Figure 7 shows a simple segmentation of the different materials obtained from the different paper target visible in Figure 2 (right). The target region has been cropped and coefficients and RGB features in the region have been used as feature for a 6-means clustering algorithm. We used the standard Matlab (R2014a) implementation using the Kmeans++ algorithm for seeding. It is evident that the intensity correction makes the material characterization obtained more meaningful. From the coefficients computed on the original images only two material are clearly separated from the others, while on the corrected images, all the paper types are effectively separated from the others by the clustering algorithm.

These results are, however, preliminary, as the final goal of our analysis is to develop a method for exploiting the planar frame and multiple spheres for a complete redesign of the PTM (or HSH) fitting. The final correction procedure will include:

- per pixel computation of a local illumination direction obtained by interpolating the directions measured on the different spheres
- direction-adapted illumination correction that, differently to the procedure proposed in Section 4, should take into account the effect of the different directions of rays illuminating the white frame

6. QUANTITATIVE ANALYSIS OF ARTWORKS

The improved detail reconstruction and material classification are, in our work, aimed at obtaining a more quantitative analysis of artworks, also to recognize materials used in paintings or to discriminate metals or stones in different conservation states. To test the possibility of discriminating different materials we created specific targets.

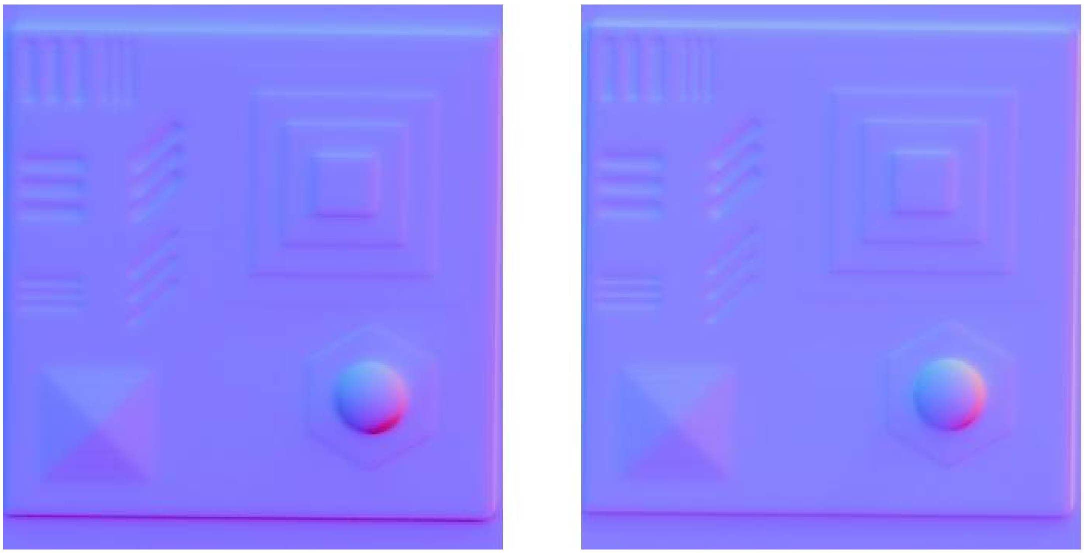

Figure 6: Left: normal estimated with RTI Viewer on the PTM reconstruction of the 3D target obtained from original images.Right: normals estimated from PTM reconstruction of intensity-corrected images. Note that contours of objects are sharperand that sphere normals show reduced artifacts.

7. CONCLUSIONS AND FUTURE WORK

We presented a preliminary experimental analysis on light intensity and directional non-uniformity in RTI acquisition. The goal is not to characterize accurately a particular setup or light source, but to provide a way to improve the classical highlight-based capture particularly widespread in the cultural heritage domain, making the results more accurate for material characterization, wit a sort of adaptive calibration procedure based on a planar frame and multiple reflective spheres.

We showed that interpolating brightness and light direction it is possible to reduce quite effectively the deviations from the uniform and "directional" light model assumed in classical PTM reconstruction. Tests performed on intensity corrected image demonstrated promising results in improving material discrimination and normal estimation. We are now working to turn this preliminary and qualitative analysis into a more quantitative one, developing a joint intensity/directionally corrected RTI fitter and measuring accurately errors on 3D reconstruction of printed targets and material classification. We plan to develop also more advanced tools allowing material characterization from highlight based RTI acquisitions also using robust fitting to segment projected shadows and highlights, as, for example proposed in [7,8].

This future activity will be also supported by a new EU funded project (Scan4Reco) where RTI will be tested for material characterization on artefacts of different types.

REFERENCES

[1] Dellepiane, M., Corsini, M., Callieri, M., & Scopigno, R. High Quality PTM Acquisition: Reflection Transformation Imaging for Large Objects. In VAST, 179-186 (2006).

[2] Earl, G., Basford, P. J., Bischoff, A. S., Bowman, A., Crowther, C., Dahl, J., ... & Kotoula, E. Reflectance transformation imaging systems for ancient documentary artefacts. BCS. (2011).

[3] Wang, O., et al. "Material classification using BRDF slices." Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on. IEEE, 2009.

[4] Jehle, M., Sommer C., and Jähne, B. "Learning of optimal illumination for material classification." Pattern Recognition. Springer Berlin Heidelberg, 2010. 563-572.

[5] Malzbender, T., Gelb, D., & Wolters, H. Polynomial texture maps. In Proceedings of the 28th annual conference on Computer graphics and interactive techniques 519-528, ACM (2001).

[6] Klaudiny, M., & Hilton, A. Error analysis of photometric stereo with colour lights. Pattern Recognition Letters, 48, 81-92 (2014).

[7] Drew, M.S, Hajari N, Hel-Or, Y and Malzbender T. Specularity and Shadow Interpolation via Robust Polynomial Texture Maps. Proc. BMVC, British Machine Vision Association, (2009)

[8] Drew, M.S., et al. "Robust estimation of surface properties and interpolation of shadow/specularity components." Image and Vision Computing 30.4 (2012): 317-331.

[9] http://culturalheritageimaging.org/What_We_Offer/Downloads/View/

[10] Arthur, D, and Vassilvitskii S.. "K-means++: The Advantages of Careful Seeding." SODA ‘07: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms. 2007, pp. 1027–1035.