

Dynamic Shading Enhancement for Reflectance Transformation Imaging

GIANPAOLO PALMA, University of Pisa and ISTI-CNR

MASSIMILIANO CORSINI, PAOLO CIGNONI, and ROBERTO SCOPIGNO, ISTI-CNR

MARK MUDGE, Cultural Heritage Imaging (CHI)

We propose a set of dynamic shading enhancement techniques for improving the perception of details, features, and overall shape characteristics from images created with Reflectance Transformation Imaging (RTI) techniques. Selection of these perceptual enhancement filters can significantly improve the user's ability to interactively inspect the content of 2D RTI media by zooming, panning, and changing the illumination direction. In particular, we present two groups of strategies for RTI image enhancement based on two main ideas: exploiting the unsharp masking methodology in the RTI-specific context; and locally optimizing the incident light direction for improved RTI image sharpness and illumination of surface features. The Result section will present a number of datasets and compare them with existing techniques.

Categories and Subject Descriptors: I.3.3 [Computer Graphics] Picture/Image Generation; I.3.6 [Computer Graphics]: Methodology and Techniques—*Interaction techniques*; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—*Color, shading, shadowing and texture*

General Terms: Algorithms, Documentation

Additional Key Words and Phrases: Visualization, cultural heritage shading enhancement, illumination, image processing

ACM Reference Format:

Palma, G., Corsini, M., Cignoni, P., Scopigno, R., and Mudge, M. 2010. Dynamic shading enhancement for reflectance transformation imaging. *ACM J. Comput. Cult. Herit.* 3, 2, Article 6 (September 2010), 20 pages.
DOI = 10.1145/1841317.1841321 <http://doi.acm.org/10.1145/1841317.1841321>

1. INTRODUCTION

Reflectance Transformation Imaging (RTI) is a digital acquisition process that captures a set of images of a subject from a single view under varying lighting conditions. RTI encodes this redundant data, the same scene sampled under many different lighting conditions, in a compact way. Given both a static object and a fixed camera view, view-dependent per-pixel reflectance functions are modeled from captured data, making it possible for the user to interactively relight the subject.

This work has been supported by a research agreement with Cultural Heritage Imaging (CHI) in the ambit of the project “Developing Advanced Technologies for the Imaging of Cultural Heritage Objects” sponsored by the US Institute of Museum and Library Services (ILMS) through a National Leadership Grant, and the Opera Primaziale of Pisa for the help and the necessary authorizations for the data acquisition.

Authors' addresses: G. Palma, Department of Computer Science, University of Pisa and Visual Computing Lab, ISTI-CNR, Pisa, Italy; M. Corsini, P. Cignoni (corresponding author), R. Scopigno, Visual Computing Lab, ISTI-CNR, Pisa, Italy; email: paolo.cignoni@isti.cnr.it; M. Mudge, Cultural Heritage Imaging (CHI), San Fransisco, CA.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2010 ACM 1556-4673/2010/09-ART6 \$10.00

DOI 10.1145/1841317.1841321 <http://doi.acm.org/10.1145/1841317.1841321>

This kind of technology has significant potential in the Cultural Heritage (CH) field, where the way light interacts with the object is important in the 3D visual observation and study of humanity's legacy. Examples of bas-reliefs, sculpture, and epigraphy are featured in this article. The characteristics of the material, reflectance behavior, and texture offer major perceptual and cognitive hints for the study of these kind of objects. Digital models encoding only the 3D shape are not able to capture all the noteworthy and interesting aspects of the artwork. In many cases the ability to interactively play with the light is often more useful than manipulation of an accurately sampled 3D shape.

Moreover, there is a wide difference in cost and precision that can be achieved with current RTI technologies compared to 3D scanning techniques: RTI techniques use inexpensive and widely available hardware (in many cases, just a digital camera and light [Mudge et al. 2006]), scale well to both large [Dellepiane et al. 2006] and very small objects [Mudge et al. 2005], and are able to easily achieve a sampling density and a precision that most current 3D scanners are unable to reach, even under optimal acquisition conditions. For those reasons, RTI techniques are widely used in the CH field for documentation tools [Mudge et al. 2008; Padfield et al. 2005] and to support detailed visual analysis [Freeth et al. 2006]. While there is an increasing interest in RTI technologies, the study of advanced visual analysis methodologies for this kind of data is in its infancy, since few enhancement techniques have been proposed.

The contribution of this article is the presentation of several shading enhancement techniques inspired by some image processing algorithms, like unsharp masking and detail shading enhancement, which we adapt to the RTI-specific domain. These techniques improve the perception of details, features, and shapes depicted in a RTI image and provide better and more flexible visual inspection capabilities with respect to traditional real-world inspection methods used by CH professionals and the existing virtual inspection tools. Real-world inspection methods present several physical limitations due to the size of the artwork and the working environment. Previous virtual inspection tools allow only the interactive relighting and the application of few methods of contrast enhancement [Malzbender et al. 2001, 2006]. We propose two groups of dynamic shading enhancement techniques based on two main ideas: to adapt the unsharp masking technique to the RTI-specific context and to create a locally adaptive *multiple-light* illumination environment. The objective is to propose a framework of methods that simplify the interaction of the user to enhance automatically or semiautomatically the most interesting details of the artwork from the point of view of the CH field.

2. RELATED WORK

Extensive literature exists regarding the definition of RTI techniques and the adoption of these applications by a wide range of fields. Reflectance transformation has been shown useful for relighting human faces, especially for digitally cloned actors used by the movie industry for special effects [Debevec et al. 2000; Wenger et al. 2005]. They have also been used for acquiring the lighting behavior and details of a variety of surfaces [Malzbender et al. 2001; Hammer et al. 2002; McGunnigle and Chantler 2001]. In this article, we will deal mainly with RTI based on the PTM framework; we refer to Malzbender et al. [2001] for an introduction to this particular kind of RTI technique. Other authors have extended this approach by defining the framework of the Bidirectional Texture Function (BTF) [Muller et al. 2005], an image-based representation that allows variation of both the lighting setup and the viewing position. We wish to emphasize that the approaches proposed in this article can be extended in an essentially straightforward way to the BTF framework.

We focus on defining the interactive rendering techniques for RTI that enable users to improve their understanding of the inspected subject. Recently, improving the readability of computer-generated images has gained significant interest and many different approaches have been proposed. On the other

hand, little attention has been devoted to investigating these issues in the RTI domain. Beyond the *specular enhancement* and *diffuse gain* techniques described in the original PTM paper [Malzbender et al. 2001], the curvature method described by Willems et al. [2005], and the surface enhancement presented in Malzbender et al. [2006], we were not aware of other RTI enhancement techniques.

2.1 Multilighting

One of the two approaches proposed in this article is based on the idea of combining, in a single view, the contributions of light coming from different directions, such that different portions of the image are illuminated in a not-physically consistent, but appealing way. This approach has been traditionally used by illustration artists to more clearly depict the form and important features of a subject. Various authors exploited this idea to build interactive, paint-like systems which combine, in a user-driven way, different portions of multiply lighted images [Akers et al. 2003; Agarwala et al. 2004]. The main idea is to create a better image by combining portions of images from a collection in which the viewpoint remains fixed and the scene static, but the lighting varies from image to image. Some authors extend this initial idea using the shading at multiple scales across the input images to generate the enhanced results in a automatic way [Fattal et al. 2007]. For each input image they compute a multiscale decomposition based on the bilateral filter and construct the enhanced output image by combining detail information from all of the input images at each scale of the decomposition. A Non-PhotoRealistic (NPR) shading model is proposed in Rusinkiewicz et al. [2006], inspired by techniques for cartographic terrain relief. The effective light position is dynamically adjusted for different areas of the surface using the local curvature of the surface along with smoothed surface normals as key elements to drive the choice of the new light position. In Toler-Franklin et al. [2007], the authors introduce RGBN images. An RGBN image stores an albedo color and a surface normal at each pixel. RGBN images are generated from several images, captured from a single camera position, of an object illuminated from different directions. This imaging method enables a user to choose from a variety of filters that are then applied to RGBN data to generate NPR renderings, such as Rusinkiewicz et al. [2006], to enhance the visibility of selected surface features. In Vergne et al. [2009], the shape depiction of 3D objects is enhanced by locally warping the environment lighting around main surface features. The method extracts the salient surface features with an analysis of the normals field in a view-dependent fashion and then it compresses or stretches locally the illumination directions with a warping function that deforms incoming light directions around salient surface features.

Concerning our approach to the previous works, we stress the fact that our technique starts from the same inspiring principle of Rusinkiewicz et al. [2006] of locally modifying the current light direction, but unlike Rusinkiewicz et al. [2006] and Vergne et al. [2009], we do not rely on the existence of a 3D model (since we have just data sampled from a single view) and our approach is based on the direct optimization of the sharpness of the resulting image instead of exploiting the local curvature. Considering Fattal et al. [2007] and Toler-Franklin et al. [2007], our approach is inherently interactive; the purpose is not to create a single image, but to provide a view-dependent interactive technique that, starting from the light direction chosen by the user, is able to create a lighting environment that maximizes the resulting sharpness and preserves a good level of brightness. In this way, the user-specified light direction is both *respected* and *enhanced*, allowing a dynamic, steerable exploration experience that simulates and improves the artifact examination overcoming the real-world physical limitations of the inspection through lighting.

2.2 Unsharp Masking

The unsharp masking approach used to improve the sharpness of a (blurred) image is the digital equivalent of a well-known darkroom analogical film processing technique [Langford 1974]. Given




Fig. 1. Projection of light vector on the image plane.

the wide adoption of this technique in image processing [Badamchizadeh and Aghagolzadeh 2004], we will focus on the previous uses of this technique in the field of 3D interactive visualization. This approach is used in Cignoni et al. [2005] to build a new set of surface normals that disclose surface discontinuities. Similarly, the unsharp masking approach is also applied in the domain of the object surface to enhance the computed shading of a given 3D scene [Ritschel et al. 2008]. In this article, we explore further the application of the unsharp masking approach in a different domain, the RTI data, following the approach of Malzbender et al. [2006]. In Malzbender et al. [2006] the unsharp masking is applied to the per-pixel normals estimated with a photometric stereo technique from the original photos that are used to make the final RTI image. We propose two extensions to this approach: using different normals which are estimated directly from the RTI image with the method proposed in Malzbender et al. [2001]; applying the unsharp masking to different data, like the luminance channel or the per-pixel coefficients of the reflectance function.

3. RTI AND PTM ENCODING

RTI images encode, given a single view of the artifact, how light coming from multiple sources is reflected by the sampled surface. A reflectance function is determined for each pixel of the image. The parameters of the reflectance function are estimated using the multiple acquired reflected light samples available for the given pixel. For example, in the case of Polynomial Texture Maps (PTM), the reflectance function is the following biquadratic polynomial

$$\begin{aligned}
 L(l_u, l_v, x, y) &= a_0(x, y) + a_1(x, y)l_u + a_2(x, y)l_v \\
 &\quad + a_3(x, y)l_u l_v + a_4(x, y)l_u^2 + a_5(x, y)l_v^2, \\
 R_f(l_u, l_v, x, y) &= L(l_u, l_v, x, y)R(x, y), \\
 G_f(l_u, l_v, x, y) &= L(l_u, l_v, x, y)G(x, y), \\
 B_f(l_u, l_v, x, y) &= L(l_u, l_v, x, y)B(x, y),
 \end{aligned} \tag{1}$$

where (l_u, l_v) is the light direction vector (normalized) projected on the image plane (see Figure 1), (a_0, \dots, a_5) are the polynomial coefficients, and (R_f, G_f, B_f) is the final color for the given pixel and the given light direction. This kind of PTM is called LRGB PTM since the “luminance” term modulates the RGB color channels. In the case of an RGB PTM each channel has its own coefficients, for a total of 18 coefficients per pixel. We refer to $L(l_u, l_v, x, y)$ as “luminance” but, more precisely, this term represents the reflectance functions with the self-shadowing effects embedded. Other basis can be substituted

for PTM to better approximate materials characterized by more complex reflectance behavior, such as gold or marble. Examples of these more sophisticated basis are spherical harmonics and hemi-spherical harmonics [Mudge et al. 2008].

In the following we refer to LRGB PTM but it is important to underline that the enhancement proposed can be easily adapted to other kinds of RTI images. Moreover, some of them can be extended also to enhance the rendering of 3D models. A brief discussion of extension will be found in the conclusions.

Finally, the term *standard rendering* used in the following examples indicates the standard PTM rendering algorithm, either in the LRGB or RGB version.

4. UNSHARP MASKING-BASED ENHANCEMENT

Unsharp masking is an image processing technique used to amplify the sharpness of a blurred image. The technique consists of enhancing the user's perception of details by extracting the fine detail information from the image and then adding this data to the original image. The fine details are calculated by subtracting a smoothed version of the original image from the original image. The smoothed/blurred version of the original image can be obtained using a gaussian or box filter. Other filters could also be employed for this purpose. Mathematically, we can formalize this operation as

$$I_E = I + k(I - I_S), \quad (2)$$

where I is the original image, I_S is the smoothed version, I_E is the enhanced image, and k is a parameter to control how much the details are amplified.

4.1 Luminance Unsharp Masking (LUM)

The first technique we propose is to apply the unsharp masking method to the $L(\cdot)$ term of the LRGB PTM (1) obtained by the current lighting conditions. In formula

$$L_E = L + k(L - L_S), \quad (3)$$

where L_S is the smoothed version of L , L_E is the enhanced one, and the parameter k controls the impact of the effect. The smoothed version, in our implementation, has been obtained by applying 5 iterations of a box filter of size 5×5 . This filtering works well for most practical cases. An optimal filter should vary this size according to the image resolution. A high number of iterations is always required to obtain a strong smoothing effect.

Since the L term, as previously stated, represents the reflectance function and includes self-shadowing effects, this way of applying the unsharp masking produces a different visual result than the standard image unsharp masking described earlier. For example, the LUM technique amplifies surface depth discontinuities. This effect is particularly evident for certain objects, such as bas-reliefs, where the perception of surface depth change can be greatly enhanced using this technique. A comparison of this technique with the standard image unsharp masking is shown in Figure 2 (see also the Result section).

4.2 Normal Unsharp Masking (NUM)

This method follows the same approach proposed in Malzbender et al. [2006] to adapt the normal enhancement technique for 3D models proposed in Cignoni et al. [2005] to RTI images. Instead to use the normals estimated from the original photos by photometric stereo techniques, we directly compute the normals from the RTI image. This permits to apply the method even to RTI images without a per-pixel saved normal.

Typically, RTI images do not contain information about surface normals, but it is possible to compute an estimation of the per-pixel normals by considering the maximum of the reflectance function. For




Fig. 2. Luminance Unsharp Masking. (Left) Standard rendering. (Center) LUM rendering. (Right) Unsharp Masking.

example, in PTM images, for each pixel the normal can be estimated considering the light direction that maximizes the L term of (1). Solving the following system

$$\frac{\partial L}{\partial u} = \frac{\partial L}{\partial v} = 0 \quad (4)$$

we obtain first the projections of surface normal (n_u, n_v)

$$n_u = \frac{a_2 a_4 - 2a_1 a_3}{4a_0 a_1 - a_2^2} \quad (5)$$

$$n_v = \frac{a_2 a_3 - 2a_0 a_4}{4a_0 a_1 - a_2^2} \quad (6)$$

and then the full normal as

$$\vec{N} = \left(n_u, n_v, \sqrt{1 - n_u^2 - n_v^2} \right). \quad (7)$$

The more Lambertian the material, the more correct the estimation of normals. Thus, specular surfaces are a clear example of the type of surface which can introduce estimation errors when we adopt such a simplified rule for the reconstruction of the normal field. The estimated normal vectors could also have another inaccuracies: since the reflectance function embeds self-shadowing effects, the maximum reflectance could be obtained from a light direction that does not coincide with the real surface normal. Our experimental results have demonstrated that, despite these approximations, in most cases the NUM approach can be used to considerably improve the readability of the image.

Therefore, the enhanced normals are obtained in the usual way

$$\vec{N}_E = \vec{N} + k(\vec{N} - \vec{N}_S), \quad (8)$$

where the smoothed normals \vec{N}_S are obtained by separately filtering each component. The smoothing filter is the same of the LUM technique.

The final rendered value is calculated considering a simple local nonphotorealistic illumination model that modulates the standard rendering

$$\begin{aligned} I &= \max(\vec{N}_E \cdot \vec{l}, 0) + k_a \\ R_f(x, y) &= R(x, y)L(l_u, l_v, x, y)I \\ G_f(x, y) &= G(x, y)L(l_u, l_v, x, y)I \\ B_f(x, y) &= B(x, y)L(l_u, l_v, x, y)I. \end{aligned} \quad (9)$$




Fig. 3. Normal Unsharp Masking (NUM). (Top-left) Standard rendering. (Top-right) Normal Unsharp Masking. (Bottom-left) Approximated normals \hat{N} . (Bottom-center) Smoothed normals \hat{N}_S . (Bottom-right) Enhanced normals \hat{N}_E .

The local illumination model is defined as the sum of a Lambertian component, the dot product between the surface normal and light vector, and a constant term k_a , which mimics an ambient term, needed to preserve the small amount of light received by the surface through indirect light and that the user can change as desired. Varying the value of k_a is possible to mitigate the artifacts caused by the nonphotorealistic illumination model used. The whole process is depicted in Figure 3.

This technique can be extended in a natural way to multichannel RTI images.

The obtained results improve the perception of the shape on the surface, especially for the objects, such as inscriptions, that have a half regular geometry.

4.3 Coefficients Unsharp Masking (CUM)

The idea of this technique is the same as LUM, to apply the unsharp masking to the basis coefficients in a component-wise way. This approach allows to enhance in a more independent way each component of the luminance modeled by the factors of the biquadratic polynomial. It is rather unclear from a theoretical point of view how to formally justify the resulting behavior as the shape of the reflectance function changes according to the “unsharped” coefficients. Such modification depends on the initial values of the coefficients, hence the visual effect is not uniform and is strongly depended on the characteristics of the material. Anyway, it has been noted in our empirical tests that this approach can give satisfying results in some cases. For these reasons, we consider CUM as an experimental technique.

The best result obtained by this technique is an enhancement of the fine details of objects that have small-scale and meso-structure features on the surface, such as grooves and bumps. Some practical examples will be shown and discussed in the Results section.

5. MULTILIGHTING DETAIL ENHANCEMENT

Multilighting detail enhancement uses different light directions to illuminate the rendered image to augment the perception of surface details. When applied in the 2D RTI domain this visualization tool works in real time and adaptively computes multiple light directions according to the portion of image under viewing and the light direction currently selected by the user. This enhancement method is able to reproduce in a very local manner what actually the archaeologists do by using a grazing light to better study and understand the artifact while preserving a good overall illumination. The advantage of this method is the possibility to create a virtual lighting environment with many lights, each of one permits to increase the local contrast of a small region of the artifact, which usually takes up some hundreds of pixel. This virtual lighting environment is not reproducible in the “real world” for several reasons: the number of lights; the very localized behavior of the light; the self-reflection effect of the object that usually decreases the sharpness of the image. The purpose of the method is to optimize an enhancement measure that maximizes the sharpness of the image and at the same time preserves the brightness. The result is a general improving of the perception of the shape and fine details present in the RTI image.

5.1 Sharpness Operator

The increased use of digital photography motivated several studies about focus measurement, with the aim of developing fast and effective auto-focus algorithms. These measurements range from simple local operators up to sophisticated and computationally expensive algorithms based on wavelet analysis. This article focuses on *sharpness operators*, simple local operators that evaluate the sharpness of an image. The idea behind the use of these operators is that maximum sharpness is achieved when the image is perfectly focused. For our purpose, these operators present two useful characteristics: they are inexpensive to compute and effective when used for contrast maximization. Some standard sharpness operators are

$$M_1 = \iint (I(x, y) - \bar{I})^2 dx dy, \quad (10)$$

$$M_2 = \iint \left| \frac{\partial I(x, y)}{\partial x} \right| + \left| \frac{\partial I(x, y)}{\partial y} \right| dx dy, \quad (11)$$

$$M_3 = \iint \left| \frac{\partial I(x, y)}{\partial x} \right|^2 + \left| \frac{\partial I(x, y)}{\partial y} \right|^2 dx dy, \quad (12)$$

$$M_4 = \iint \left| \frac{\partial^2 I(x, y)}{\partial x^2} \right|^2 + \left| \frac{\partial^2 I(x, y)}{\partial y^2} \right|^2 dx dy, \quad (13)$$

$$M_5 = \iint (\nabla^2 I(x, y))^2 dx dy, \quad (14)$$

where M_1 is the image variance, M_2 is the L^1 -norm of the image gradient, M_3 is the L^2 -norm of the image gradient, M_4 is the L^1 -norm of the second derivatives, and M_5 is the energy of the image Laplacian.

In our implementation of the proposed algorithm we tested all those operators that are based on the first and on the second derivatives, in particular the M_2 , M_3 , and M_5 operators. The image gradient is estimated with a Sobel filter, while the Laplacian is calculated using an image Laplace operator [Gonzalez and Woods 2008].

5.2 Multilighting Dynamic Enhancement

Typically, the visualization tools for RTI images allow the user to specify interactively the light direction. Zoom and pan operations are generally available to visualize high-resolution images, usually together with a multiresolution encoding. To speed up the rendering process a mip-mapping approach is used. The original RTI image is scaled at different levels of resolution. These levels are used according to the zoom factor of the browser.

In the following we indicate the user-selected light direction with l_0 and the section of the image currently rendered with I^* . The dynamic multilighting detail enhancement algorithm can be summarized by the following steps.

- (1) The image under viewing (I^*) is subdivided into $N \times M$ square tiles.
- (2) For each tile (T) l_0 is perturbed, the light direction (l') which maximizes the detail enhancement is chosen as the light direction for that tile.
- (3) The grid of NM light directions obtained in this way are then made more uniform by applying a smoothing filter component-wise.
- (4) Finally, per-tile light directions are converted into per-pixel light directions through bilinear interpolation.

The dimensions of the tiles grid $M \times N$ depend on the tile size (in pixel) chosen by the user, but they are almost independent of the browser resolution. In fact, for the evaluation we use the level of the mip-mapping with the resolution closer to the current view resolution while taking into account the zoom factor. In this way the number of tiles in any view is more or less constant and the lighting configuration is recalculated during a zoom-in operation in order to reveal better fine details. The Results section will show more about this feature.

The step 2 of the algorithm requires to solve the optimization problem

$$l' = \arg \max_{l \in \mathcal{L}(l_0)} \mathcal{E}(T, l), \quad (15)$$

where $\mathcal{L}(l_0)$ is the set of perturbed light directions computed from l_0 and \mathcal{E} is the enhancement measure described as follows.

The set of lights $\mathcal{L}(l_0)$ is computed by perturbing l_0 . We use two different methods to obtain the set of light direction: *anisotropic* sampling and *isotropic* sampling. With the anisotropic sampling the l_0 set is generated such that the amount of perturbation decrease for the light directions that are almost tangent to the viewing plane (see Figure 4). This sampling strategy allows careful treatment of the case starting from a grazing light direction. In this condition, even small perturbations can result in high visual changes. Hence, in this case, we prefer to perturb the light direction with rotations along the viewing axis rather than with rotations perpendicular to it.

In practice, we take advantage of the empirical knowledge that our RTI images will probably represent something that has an “interesting” behavior under grazing light and therefore we develop a targeted optimization strategy that takes this critical configuration into account. Additionally, the isotropic sampling permits the light set to uniformly spread in a cone of direction with axis l_0 that keeps the same size independently from the direction of l_0 . A comparison of the two methods is shown in Figure 6. As can be seen, the anisotropic sampling preserves more of the dark-and-light parts of the image with respect to the standard rendering illuminated with grazing light while avoiding visual artifacts. The isotropic sampling tends to illuminate the image more uniformly but it can potentially generate artifacts if abrupt lighting changes occur.




Fig. 4. Anisotropic sampling of light directions. Each ellipse shows the direction sampled assuming that l_0 is the ellipse center.




Fig. 5. Light fields computed on-the-fly. Each vector represents the light direction used for that tile.

The enhancement measure is defined as

$$\mathcal{E}(T, l) = \alpha \mathcal{S}(T(l)) + (1 - \alpha) \mathcal{Y}(T(l)), \quad (16)$$

where $\mathcal{S}(\cdot)$ is the sharpness of the tile T evaluated using one of the previously described sharpness operators, $\mathcal{Y}(\cdot)$ is a measure of the total brightness of the tile, and α is a tuning parameter that



Fig. 6. Dynamic multilighting detail enhancement ($\alpha = 0.70$).

controls the sharpness/brightness ratio. The tile's brightness is evaluated by converting the RGB color component to YUV color space and by summing up the luminance for each pixel. The two measures are scaled according to the maximum of their values in order to have compatible values. The introduction of the brightness term in the enhancement metric $\mathcal{E}(T, l)$ is very important because the sharpness maximization alone tends to exaggerate black-white contrast, making the final image too dark in some




Fig. 7. Multiresolution lighting constraint. The light direction at the top level of subdivision influences the light directions at the successive levels of subdivision.



Fig. 8. Static multilighting detail enhancement with anisotropic sampling. (Top) Standard rendering. (Bottom-left) Enhancement metric with $\alpha = 0.7$. (Bottom-right) Enhancement metric with $\alpha = 0.4$. It is possible to notice the left image has a high contrast but it is too dark due to the prevalence of contrast over brightness ($\alpha = 0.7$). The right image is well balanced (good contrast, good overall illumination).

cases. Thanks to the brightness term we are able to generate high-contrast images while preserving a good global illumination over the entire image (see Figure 8).

After the lighting vectors are calculated for each tile, a smoothing filter is applied during step 3 of the algorithm. Though the parameter of the filter can be chosen by the user, in many cases we apply a box filter of size 3×3 between 5 to 7 times to produce a strong smoothing effect. This filtering step is necessary since the light directions in adjacent tiles may differ considerably and can produce visible artifacts.

During the final rendering (step 4) a per-pixel light direction is used to avoid visual differences across the tiles. For each pixel the light direction is obtained by bilinear interpolation of the light directions of the four tiles adjacent to it. An example of the light field computed as described is shown in Figure 5.

One of the main advantages and innovative aspects of this enhancement method is its view-dependent nature. A zoom operation reveals more details than an overview of a big image. Since the lighting field is calculated, for each frame, in a screen-space tile-based way, the scale of the details

which are enhanced is automatically adjusted. As a side effect, the variation of light direction on the surface of the object is somewhat bounded in screen space: the more the user zooms in on a detail the more the light directions bend to the local features. This is very important for very large images. When looking at the whole image, the light directions could vary too fast, leading to significantly unpleasant visual effects.

5.3 Multilighting Static Enhancement

The dynamic multilighting enhancement can be easily modified to produce an automatic high-contrast, well-illuminated image for stand-alone presentations, high-quality printing, or similar purposes. We call this version of the algorithm *static multilighting enhancement*. The static version of the algorithm differs from the dynamic one in two aspects: the lighting field is calculated on the entire image instead of on I^* and all the hemisphere are sampled during the detail enhancement optimization step. In this way, hundreds of local lighting configurations are tested and the best one according to the enhancement metric (Eq. (16)) selected for the final rendering.

This modification addresses the problem that great lighting variations across the image can produce visual artifacts. We introduce some constraints in the light field generation to overcome this problem and use all the possible lighting directions during the image enhancement and rendering phase. The constrained light field is generated by means of multiresolution framework (see Figure 7). At the first iteration the entire image is considered and all the possible light directions are evaluated (256 directions in our tests) according to the \mathcal{E} metric (16) to estimate the optimal light direction (l'_0). At the second iteration the image is subdivided into four tiles and, for each tile, the optimal light is computed starting from l'_0 . In general, after the first iteration the optimal light direction for each tile is chosen among a set of light vectors computed by perturbing the optimal direction of the father tile with the sampling method described for the dynamic enhancement; either anisotropic or isotropic sampling. The process continues in the same way until we reach a predefined level of subdivision.

In order to reduce the computational cost of the algorithm the initial iterations work on a subsampled version of the image.

6. RESULTS AND DISCUSSION

In this section we show some images generated with the proposed techniques, while describing and discussing the visual effects obtained. In general we have received very positive opinions by the users, especially CH professionals, about these new methods and their results. Anyway, it should be important to schedule as a future work a series of perceptual tests to obtain a more objective evaluation. On the other hand, we can state that the proposed multilighting methods allow to produce a sharper image because we maximize the global sharpness term of the enhancement measure of Eq. (16).

Some of the PTMs used in the experiments are provided with the PTM software distributed by the HP Laboratories, for example, the towel cookie board (Figure 2) and the cuneiform tablet (Figure 3). The other PTMs depict a sarcophagus in the monumental cemetery in Pisa (Figure 9), a bas-relief in exposition at the Museo of the Opera del Duomo in Pisa (Figure 5), and a high-relief in gilded wood representing a kiss between Corsica and Elba islands, from the Isola d'Elba museum, Italy (Figure 8). Information about the rendering times are given at the end of this section.

The images rendered with LUM are very effective, the perception of the image details is in general improved with respect to the standard rendering (see Figures 2 and 9). The main differences with the standard unsharp masking is that LUM enhances the luminance difference and depth discontinuities of the image, for example, the spongy material on the background of Figure 2, while unsharp masking deals with color variations and does not take in account the subtend geometric features. This effect is



Fig. 9. Classic and luminance unsharp masking.

particularly evident in the wood's grain in Figure 2. For the same reason, LUM produces images that are “cleaner” than image unsharp masking in the general case (see Figure 9).

Concerning the NUM technique, the experimental results are encouraging even if, in some cases, a few visual artifacts are produced. These artifacts are generated when the normals are not well estimated due to the approximations mentioned before. The final effect is the presence of clusters of completely black pixel in areas where the angle between the normal and the light direction is greater than $\pi/2$. To reduce these artifacts the user can increase the term k_a of the nonphotorealistic illumination model used in Eq. (10). An example of these artifacts is shown in Figure 11. This technique still remains worthwhile to analyze, for example, inscriptions which are difficult to read otherwise (see Figure 10).

Even if CUM is an experimental technique and its results need further study to be well understood, in most cases this technique provides an interesting enhancement of the fine details, like small grooves and bumps, as shown in Figures 12 and 10.

The multilighting detail enhancement can produce different visual effects thanks to the flexibility of the algorithm. The parameter α allows the user to tune the enhancement metric to improve the contrast or the global brightness of the image. Figure 13 shows an example where the parameter α permits an increase in the depth perception and the sharpness of the image and, at the same time, maintains a good global brightness. Figure 8 shows the effects of the enhancement measure (16) varying the parameter α in the static multilighting detail enhancement.

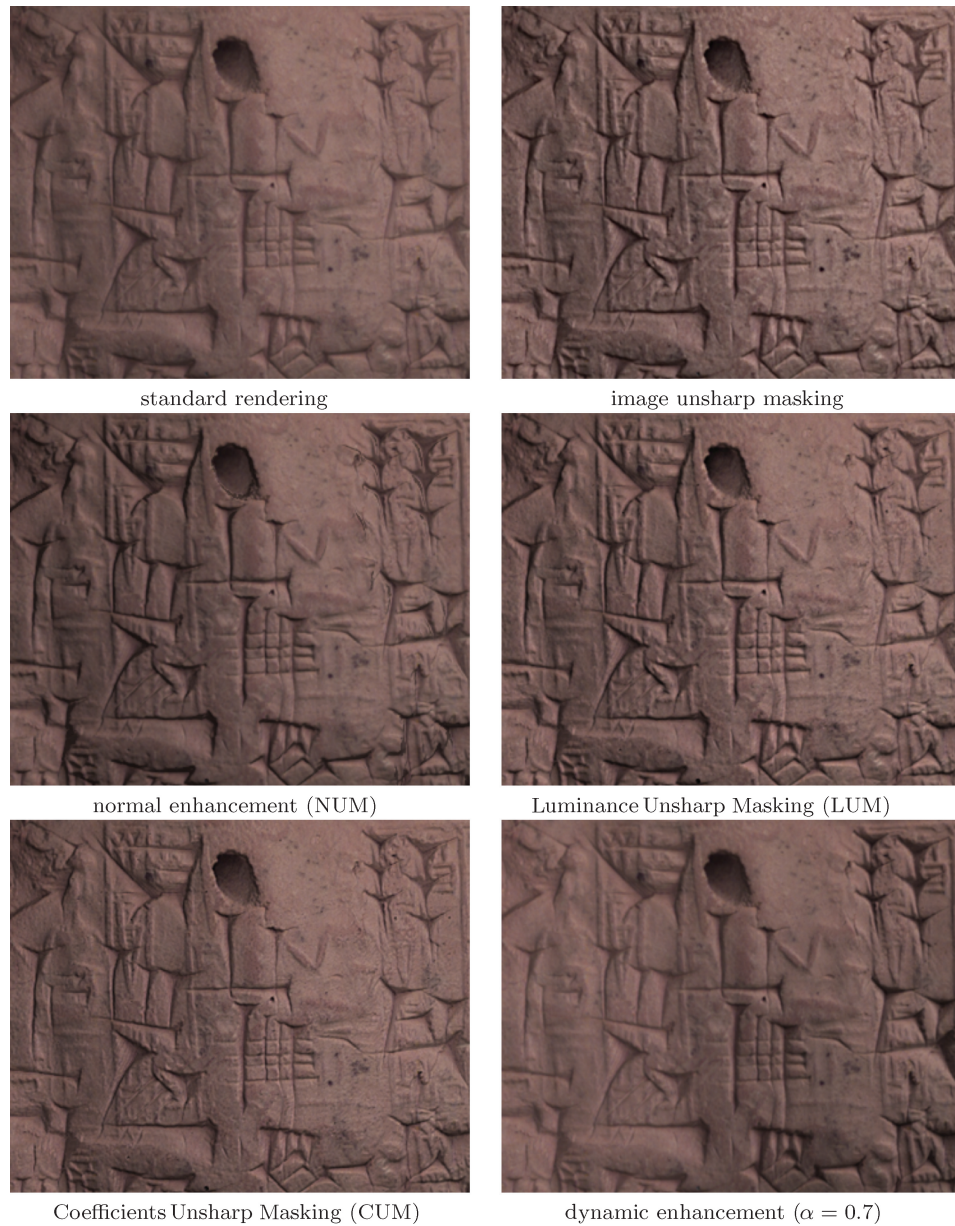


Fig. 10. Techniques comparison. Detail of the cuneiform tablet.

Regarding the use of different sharpness operators, we have tested the M_2 , M_3 , and M_5 operators (14). In general, the Laplacian energy gives the best results, most probably because the second derivative is more accurate in edge and fine detail extraction. The Laplacian energy is influenced by the image content, hence the final results present greater light variations. The L^1 -norm and L^2 -norm of the image gradient generate images with more uniform illumination. The L^2 -norm can considerably



Fig. 11. Visual artifacts in the NUM technique. (Top-left) Standard rendering. (Top-right) NUM rendering with $k = 1.0$ and $k_\alpha = 0.5$. (Bottom-left) NUM rendering with $k = 4.0$ and $k_\alpha = 0.5$. (Bottom-right) NUM rendering with $k = 4.0$ and $k_\alpha = 1.5$.

amplify the details in some cases. An example of the use of different sharpness operators is shown in Figure 15 where a detail of the Corsica-Elba PTM is depicted.

Another important parameter of multilighting detail enhancement is light sampling. The main difference among the two methods proposed is shown in Figure 6. With anisotropic sampling the lighting tends to be grazing, because we constrain the light direction to rotate along the viewing axis; while with the isotropic sampling it becomes more perpendicular to the image surface. Furthermore, the two methods have opposite disadvantages: anisotropic sampling can generate excessively dark images in some parts; isotropic sampling can generate images that are too bright. A possible solution is to decrease the α value for the former and to increase it for the later.

Finally, we will discuss performance. All the techniques presented, with the exception of static detail enhancement, work at interactive or near-interactive rates, despite the fact that initial implementation was limited to the CPU. The rendering times on a commodity PC are summarized in Table I. The image used in the test has a resolution of 2516×1646 and takes about 106Mb of memory, while the size of the view in the visualization tool is 800×800 pixel with zoom factor 1:1. The rendering time for the static



Fig. 12. (Left) Standard rendering. (Right) Coefficients Unsharp Masking (CUM).



Fig. 13. (Left) Standard rendering. (Right) dynamic multilighting detail enhancement with anisotropic sampling ($\alpha = 0.8$).

multilighting technique is high because it uses the whole image to generate the final static image. We remember that this technique is not interactive but designed for presentation purposes. It is important to emphasize that all the techniques proposed can be implemented in GPU, resulting in an expected great increase of performances. This is an improvement we have planned to do in our future work.

7. CONCLUSION

We have presented two groups of dynamic shading enhancement techniques for improving the perception of details, features, and shape within images created with RTI techniques. The first group of strategies is based on the idea of exploiting the unsharp masking methodology in the RTI-specific context, while the second approach focuses on the task of locally optimizing the light direction for



Fig. 14. Dynamic detail enhancement with anisotropic sampling. (Top-left) Standard rendering. (Top-right) Dynamic detail enhancement with 32 tiles and 20° of light perturbation ($\alpha = 0.7$). (Bottom-left) Standard rendering. (Bottom-right) Dynamic detail enhancement with 32 tiles and 20° of light perturbation ($\alpha = 0.65$). Changes in scale are automatically achieved thanks to the view-dependent nature of the algorithm.

Table I. Rendering Times (Intel Core Duo 1800Ghz, 2GB Ram)

Enhancement Techniques	Rendering Times (sec)
Standard	0.051
LUM	0.168
NUM	0.182
Dynamic Multi-Lighting	0.895
Static Multi-Lighting	43.57

improving the sharpness and brightness of the resulting image. Note that most of the proposed techniques are not closely tied with the specific PTM encoding, but are applicable to different types of RTI basis function and to more complex image-based approaches like BTF.

It is interesting to note that even if the application of unsharp masking to a different domain is not a novel idea, the application of unsharp masking to RTI images enables several visual enhancement effects that enrich remote visual inspection of our cultural heritage.




Fig. 15. Sharpness operator comparison in the static detail enhancement with $\alpha = 0.65$. (Top-left) Standard rendering. (Top-right) Energy of Laplacian. (Bottom-left) L^1 -norm of the image gradient (Bottom-right) L^2 -norm of the image gradient.

The strongest and most novel RTI technique proposed is multilighting enhancement. This algorithm can be extended naturally to the visualization of 3D models. This extension is possible by applying the algorithm as described directly in image space.

The algorithms proposed here have been included in a novel tool for RTI image visualization, called *RTIViewer*, which includes on-demand loading of multiresolution RTI images from Web and several visualization modes. The tool is the result of a joined research of Cultural Heritage Imaging (CHI) and Visual Computing Laboratory (ISTI-CNR) and will be released soon. For further information see [C-H-I 2010]. Even if we have received positive feedback by the archaeologists that have used the tool and the new enhancement methods, as future work it would be interesting to conduct several perceptual tests to obtain a more objective evaluation of the improvement of these new techniques.

REFERENCES

- AGARWALA, A., DONTCHEVA, M., AGRAWALA, M., DRUCKER, S., COLBURN, A., CURLESS, B., SALESIN, D., AND COHEN, M. 2004. Interactive digital photomontage. *ACM Trans. Graph.* 23, 3, 294–302.
- AKERS, D., LOSASSO, F., KLINGNER, J., AGRAWALA, M., RICK, J., AND HANRAHAN, P. 2003. Conveying shape and features with image-based relighting. In *Proceedings of the 14th IEEE Visualization Conference (VIS'03)*, 349–354.
- BADAMCHIZADEH, M. AND AGHAGOLZADEH, A. 2004. Comparative study of unsharp masking methods for image enhancement. In *Proceedings of the 3rd International Conference on Image and Graphics*. 27–30.
- C-H-I. 2010. Cultural Heritage Imaging - RTI resources. <http://www.c-h-i.org/learn/index.html>.
- CIGNONI, P., SCOPIGNO, R., AND TARINI, M. 2005. A simple normal enhancement technique for interactive non-photorealistic renderings. *Comput. Graph.* 29, 1, 125–133.

- DEBEVEC, P., HAWKINS, T., TCHOU, C., DUKER, H.-P., SAROKIN, W., AND SAGAR, M. 2000. Acquiring the reflectance field of a human face. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'00)*. ACM Press, New York, 145–156.
- DELLEPIANE, M., CORSINI, M., CALLIERI, M., AND SCOPIGNO, R. 2006. High quality ptm acquisition: Reflection transformation imaging for large objects. In *Proceedings of the 7th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST'05)*. M. Ioannides, D. Arnold, F. Niccolucci, and K. Mania, Eds., Eurographics Association, 179–186.
- FATTAL, R., AGRAWALA, M., AND RUSINKIEWICZ, S. 2007. Multiscale shape and detail enhancement from multi-light image collections. *ACM Trans. Graph.* 26, 3.
- FREETH, T., BITSAKIS, Y., MOUSSAS, X., SEIRADAKIS, J., TSELIKAS, A., MANGOU, H., ZAFEIROPOULOU, M., HADLAND, R., BATE, D., RAMSEY, A., ET AL. 2006. Decoding the ancient greek astronomical calculator known as the antikythera mechanism. *Nature* 444, 7119, 587–591.
- GONZALEZ, R. AND WOODS, R. 2008. *Digital Image Processing*, 3rd ed. Prentice Hall.
- HAMMER, Ø., BENGTON, S., MALZBENDER, T., AND GELB, D. 2002. Imaging fossils using reflectance transformation and interactive manipulation of virtual light sources. *Palaeontologia Electronica* 5, 1, 9.
- LANGFORD, M. 1974. *Advanced Photography: A Grammar of Techniques*. Focal Press.
- MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial texture maps. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH'03)*. ACM, 519–528.
- MALZBENDER, T., WILBURN, B., GELB, D., AND AMBRISCO, B. 2006. Surface enhancement using real-time photometric stereo and reflectance transformation. In *Proceedings of the Eurographics Workshop/Symposium on Rendering*. T. Akenine-Möller and W. Heidrich, Eds., Eurographics Association, 245–250.
- MCGUNNIGLE, G. AND CHANTLER, M. 2001. Recovery of fingerprints using photometric stereo. In *Proceedings of the Irish Machine Vision and Image Processing Conference (IMVIP'01)*. 192–199.
- MUDGE, M., MALZBENDER, T., CHALMERS, A., SCOPIGNO, R., DAVIS, J., WANG, O., GUNAWARDANE, P., ASHLEY, M., DOERR, M., PROENCA, A., AND BARBOSA, J. 2008. Image-based empirical information acquisition, scientific reliability, and long-term digital preservation for the natural sciences and cultural heritage. In *(Eurographics'08) Tutorials*, M. Roussou and J. Leigh, Eds., Eurographics Association.
- MUDGE, M., MALZBENDER, T., SCHROER, C., AND LUM, M. 2006. New reflection transformation imaging methods for rock art and multiple-viewpoint display. In *Proceedings of the International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage (VAST'06)*. M. Ioannides, D. Arnold, F. Niccolucci, and K. Mania, Eds., Eurographics Association, 195–202.
- MUDGE, M., VOUTAZ, J.-P., SCHROER, C., AND LUM, M. 2005. Reflection transformation imaging and virtual representations of coins from the hospice of the grand st. bernard. In *Proceedings of the 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage*. M. Mudge, N. Ryan, and R. Scopigno, Eds., Eurographics Association. 29–39.
- MULLER, G., MESETH, J., SATTLER, M., SARLETTE, R., AND KLEIN, R. 2005. Acquisition, synthesis, and rendering of bidirectional texture functions. *Comput. Graph. Forum* 24, 1, 83–109.
- PADFIELD, J., SAUNDERS, D., AND MALZBENDER, T. 2005. Polynomial texture mapping: A new tool for examining the surface of paintings. In *ICOM Committee for Conservation Triennial Meeting*. Eurographics Association, 504–510.
- RITSCHER, T., SMITH, K., IHRKE, M., GROSCH, T., MYRSKOWSKI, K., AND SEIDEL, H.-P. 2008. 3d unsharp masking for scene coherent enhancement. In *ACM (SIGGRAPH'08) Papers*. ACM, New York, 1–8.
- RUSINKIEWICZ, S., BURNS, M., AND DECARLO, D. 2006. Exaggerated shading for depicting shape and detail. *ACM Trans. Graph.* 25, 3, 1199–1205.
- TOLER-FRANKLIN, C., FINKELSTEIN, A., AND RUSINKIEWICZ, S. 2007. Illustration of complex real-world objects using images with normals. In *Proceedings of the 5th International Symposium on Non-Photorealistic Animation and Rendering*. B. Gooch, M. Agrawala, and O. Deussen, Eds., ACM, 111–119.
- VERGNE, R., PACANOWSKI, R., BARLA, P., GRANIER, X., AND SCHLICK, C. 2009. Light warping for enhanced surface depiction. *ACM Trans. Graph.* 28, 3.
- WENGER, A., GARDNER, A., TCHOU, C., UNGER, J., HAWKINS, T., AND DEBEVEC, P. 2005. Performance relighting and reflectance transformation with time-multiplexed illumination. In *ACM Trans. Graph.* 24, 3, 756–764.
- WILLEMS, G., VERBIEST, F., MOREAU, W., HAMEEUW, H., VAN LERBERGHE, K., AND VAN GOOL, L. 2005. Easy and cost-effective cuneiform digitizing. In *Proceedings of the 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage*. M. Mudge, N. Ryan, and R. Scopigno, Eds., Eurographics Association, 73–80.

Received March 2009; revised May 2010; accepted July 2010