



HAL
open science

Radiance Scaling for Versatile Surface Enhancement

Romain Vergne, Romain Pacanowski, Pascal Barla, Xavier Granier,
Christophe Schlick

► **To cite this version:**

Romain Vergne, Romain Pacanowski, Pascal Barla, Xavier Granier, Christophe Schlick. Radiance Scaling for Versatile Surface Enhancement. I3D '10: Proc. symposium on Interactive 3D graphics and games, Feb 2010, Boston, United States. 10.1145/1730804.1730827 . inria-00449828

HAL Id: inria-00449828

<https://inria.hal.science/inria-00449828v1>

Submitted on 9 Jun 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Radiance Scaling for Versatile Surface Enhancement




Figure 1: Our novel Radiance Scaling technique enhances the depiction of surface shape under arbitrary illumination, with various materials, and in a wide range of rendering settings. In the left pair of images, we illustrate how surface features are enhanced mainly through enhancement of the specular shading term. Whereas on the right pair of images, we show the efficiency of our method on an approximation of a refractive material. Observe how various surface details are enhanced in both cases: around the eyes, inside the ear, and on the nose.

Abstract

We present a novel technique called Radiance Scaling for the depiction of surface shape through shading. It adjusts reflected light intensities in a way dependent on both surface curvature and material characteristics. As a result, diffuse shading or highlight variations become correlated to surface feature variations, enhancing surface concavities and convexities. This approach is more versatile compared to previous methods. First, it produces satisfying results with any kind of material: we demonstrate results obtained with Phong and Ashikmin BRDFs, Cartoon shading, sub-Lambertian materials, and perfectly reflective or refractive objects. Second, it imposes no restriction on lighting environment: it does not require a dense sampling of lighting directions and works even with a single light. Third, it makes it possible to enhance surface shape through the use of precomputed radiance data such as Ambient Occlusion, Prefiltered Environment Maps or Lit Spheres. Our novel approach works in real-time on modern graphics hardware and is faster than previous techniques.

Keywords: Expressive rendering, NPR, Shape depiction.

1 Introduction

The depiction of object shape has been a subject of increased interest in the Computer Graphics community since the work of Saito and Takahashi [1990]. Inspired by their pioneering approach, many rendering techniques have focused on finding an appropriate set of lines to depict object shape. In contrast to line-based approaches, other techniques depict object shape through shading. Maybe the most widely used of these is *Ambient Occlusion* [Pharr and Green 2004], which measures the occlusion of nearby geometry. Both types of techniques make drastic choices for the type of material, illumination and style used to depict an object: line-based approaches often ignore material and illumination and depict mainly sharp surface features, whereas occlusion-based techniques convey

deep cavities for diffuse objects under ambient illumination.

More versatile shape enhancement techniques are required to accommodate the needs of modern Computer Graphics applications. They should work with realistic as well as stylized rendering to adapt to the look-and-feel of a particular movie or video game production. A wide variety of materials should be taken into account, such as diffuse, glossy and transparent materials, with specific controls for each material component. A satisfying method should work for various illumination settings ranging from complex illumination for movie production, to simple or even precomputed illumination for video games. On top of these requirements, enhancement methods should be fast enough to be incorporated in interactive applications or to provide instant feedback for previewing.

This versatility has been recently tackled by techniques that either modify the final evaluation of reflected radiance as in *3D Unsharp masking* [Ritschel et al. 2008], or modify it for each incoming light direction as in *Light Warping* [Vergne et al. 2009]. These techniques have shown compelling enhancement abilities without relying on any particular style, material or illumination constraint. Unfortunately, as detailed in Section 2, these methods provide at best a partial control on the enhancement process and produce unsatisfying results or even artifacts for specific choices of material or illumination. Moreover, both methods are dependent on scene complexity: 3D Unsharp Masking performances slow down with an increasing number of visible vertices, whereas Light Warping requires a dense sampling of the environment illumination, with a non-negligible overhead per light ray.

The main contribution of this paper is to present a technique to depict shape through shading that combines the advantages of 3D Unsharp Masking and Light Warping while providing a more versatile and faster solution. The key idea is to adjust reflected light intensities in a way that depends on both surface curvature and material characteristics, as explained in Section 3. As with 3D Unsharp Masking, enhancement is performed by introducing variations in reflected light intensity, an approach that works for any kind of illu-

mination. However, this is not performed indiscriminately at every surface point and for the outgoing radiance only, but in a curvature-dependent manner and for each incoming light direction as in Light Warping. The main tool to achieve this enhancement is a novel scaling function presented in Section 4. In addition, Radiance Scaling takes material characteristics into account, which not only allows users to control accurately the enhancement per material component, but also makes the method easy to adapt to different rendering scenarios as shown in Section 5. Comparisons with related techniques and directions for future work are given in Section 6.

2 Previous work

Most of the work done for the depiction of shape in Computer Graphics concerns line-based rendering techniques. Since the seminal work of Saito and Takahashi [1990], many novel methods (e.g., [Nienhaus and Döllner 2004; Ohtake et al. 2004; DeCarlo et al. 2003; Judd et al. 2007; Lee et al. 2007; Goodwin et al. 2007; Kolomenkin et al. 2008; Zhang et al. 2009]) have been proposed. Most of these techniques focus on depicting shape features directly, and thus make relatively little use of material or illumination information, with the notable exception of Lee et al. [2007].

A number of shading-based approaches have also shown interesting abilities for shape depiction. The most widely used of these techniques is *Ambient Occlusion* [Pharr and Green 2004], which measures the occlusion of nearby geometry. The method rather tends to depict deep cavities, whereas shallow (yet salient) surface details are often missed or even smoothed out. Moreover, enhancement only occurs *implicitly* (there is no direct control over the shading features to depict), and the method is limited to diffuse materials and ambient lighting. It is also related to *Accessibility shading* techniques (e.g., [Miller 1994]), which conveys information about concavities of a 3D object.

The recent *3D Unsharp Masking* technique of Ritshel et al. [2008] addresses limitations on the type of material or illumination. It consists in applying the Cornsweet Illusion effect to outgoing radiance on an object surface. The approach provides interesting enhancement not only with diffuse materials, but also with glossy objects, shadows and textures. However, the method is applied indiscriminately to all these effects, and thus enhances surface features only implicitly, when radiance happens to be correlated with surface shape. Moreover, it produces artifacts when applied to glossy objects: material appearance is then strongly altered and objects tend to look sharper than they really are. Hence, the method is likely to create noticeable artifacts when applied to highly reflective or refractive materials as well.

In this paper, we rather seek a technique that enhances object shape *explicitly*, with intuitive controls for the user. Previous methods [Kindlmann et al. 2003; Cignoni et al. 2005; Rusinkiewicz et al. 2006; Vergne et al. 2008; Vergne et al. 2009] differ in the geometric features they enhance and on the constraints they put on materials, illumination or style. For instance, *Exaggerated Shading* [Rusinkiewicz et al. 2006] makes use of normals at multiple scales to define surface relief and relies on a Half-Lambertian to depict relief at grazing angles. The most recent and general of these techniques is *Light Warping* [Vergne et al. 2009]. It makes use of a view-centered curvature tensor to define surface features, which are then enhanced by locally stretching or compressing reflected light patterns around the view direction. Although this technique puts no constraint on the choice of material or illumination, its effectiveness decreases with lighting environments that do not exhibit natural statistics. It also requires a dense sampling of illumination, and is thus not adapted to simplified lighting such as found in video games, or to the use of precomputed radiance methods. Moreover,

highly reflective or refractive materials produce complex warped patterns that tend to make rendering less legible.

3 Overview

The key observation of this paper is that *explicitly* correlating reflected lighting variations to surface feature variations leads to an improved depiction of object shape. For example, consider a highlight reflected off a glossy object; by increasing reflected light intensity in convex regions and decreasing it in concave ones, the highlight looks as if it is attracted toward convexities and repelled from concavities (see Figure 1-left). Such an adjustment improves the distinction between concave and convex surface features, and does not only take surface features into account, but also material characteristics. Indeed, reflected light intensity has an altogether different distribution across the surface depending on whether the material is glossy or diffuse for instance.

The main idea of Radiance Scaling is thus to adjust reflected light intensity per incoming light direction in a way that depends on both surface curvature and material characteristics. Formally, we rewrite the reflected radiance equation as follows:

$$L'(\mathbf{p} \rightarrow \mathbf{e}) = \int_{\Omega} \rho(\mathbf{e}, \boldsymbol{\ell})(\mathbf{n} \cdot \boldsymbol{\ell}) \sigma(\mathbf{p}, \mathbf{e}, \boldsymbol{\ell}) L(\mathbf{p} \leftarrow \boldsymbol{\ell}) d\boldsymbol{\ell} \quad (1)$$

where L' is the enhanced radiance, \mathbf{p} is a surface point, \mathbf{e} is the direction toward the eye, \mathbf{n} is the surface normal at \mathbf{p} , Ω is the hemisphere of directions around \mathbf{n} , $\boldsymbol{\ell}$ is a light direction, ρ is the material BRDF, σ is a scaling function and L is the incoming radiance.

The scaling function is a short notation for $\sigma_{\alpha, \gamma}(\kappa(\mathbf{p}), \delta(\mathbf{e}, \boldsymbol{\ell}))$. The curvature mapping function $\kappa(\mathbf{p}) : \mathbb{R}^3 \rightarrow [-1, 1]$ computes normalized curvature values, where -1 corresponds to maximum concavities, 0 to planar regions and 1 to maximum convexities. We call $\delta(\mathbf{e}, \boldsymbol{\ell}) : \Omega^2 \rightarrow [0, 1]$ the reflectance mapping function. It computes normalized values, where 0 corresponds to minimum reflected intensity, and 1 to maximum reflected intensity. Intuitively, it helps identify the light direction that contributes the most to reflected light intensity.

We describe the formula for the scaling function and the choice of curvature mapping function in Section 4. We then show how Radiance Scaling is easily adapted to various BRDF and illumination scenarios by a proper choice of reflectance mapping function in Section 5.

4 Scaling function

The scaling term in Equation 1 is a function of two variables: a normalized curvature and a normalized reflectance. Both variables are themselves functions, but for clarity of notation, we use κ and δ in the following. We require the scaling function to be monotonic, so that no new shading extremum is created after scaling. Another requirement is that for planar surface regions, the function must have no influence on reflected lighting. The following function fulfills these requirements (see Figure 2):

$$\sigma_{\alpha, \gamma}(\kappa, \delta) = \frac{\alpha \gamma e^{\kappa} + \delta(1 - \alpha(1 + \gamma e^{\kappa}))}{(\alpha + \delta(\gamma e^{\kappa} - \alpha(1 + \gamma e^{\kappa})))} \quad (2)$$

where $\alpha \in (0, 1)$ controls the location of the scaling-invariant point of σ and $\gamma \in [0, \infty)$ is the scaling magnitude. The scaling-invariant point is a handy parameter to control how variations in shading depict surface feature variations. For convex features, reflected lighting intensities above α are brightened and those below α are darkened. For concave features, the opposite effect is obtained. This is illustrated in Figure 3.




Figure 2: Two plots of a set of scaling functions with different scaling-invariant points (left: $\alpha = 0.2$; right: $\alpha = 0.8$), and using increasing curvatures $\kappa = \{-1, -1/2, 0, 1/2, 1\}$.




Figure 3: The effect of scaling parameters. Left: no scaling ($\gamma = 0$). Middle: scaling with a low scaling-invariant point ($\alpha = 0.2$): convexities are mostly brightened. Right: scaling with a high scaling-invariant point ($\alpha = 0.8$): convexities are brightened in the direction of the light source, but darkened away from it.

5 Rendering scenarios

We now explain how the choice of reflectance mapping function δ permits the enhancement of surface features in a variety of rendering scenarios. Reported performances have been measured at a 800×600 resolution using a NVIDIA Geforce 8800 GTX.

5.1 Simple lighting with Phong shading model

In interactive applications such as video games, it is common to make use of simple shading models such as Phong shading, with a restricted number of light sources. Radiance Scaling allows users to control each term of Phong’s shading model independently, as explained in the following.

With a single light source and Phong shading, Equation 1 becomes

$$L'(\mathbf{p} \rightarrow \mathbf{e}) = \sum_j \rho_j(\mathbf{e}, \boldsymbol{\ell}_0) \sigma_j(\mathbf{p}, \mathbf{e}, \boldsymbol{\ell}_0) L_j(\boldsymbol{\ell}_0)$$

where $j \in \{a, d, s\}$ iterates over the ambient, diffuse and specular components of Phong’s shading model and $\boldsymbol{\ell}_0$ is the light source direction at point \mathbf{p} . For each component, L_j corresponds to light intensity (L_a is a constant). The ambient, diffuse and specular components are given by $\rho_a = 1$, $\rho_d(\boldsymbol{\ell}_0) = (\mathbf{n} \cdot \boldsymbol{\ell}_0)$ and $\rho_s(\mathbf{e}, \boldsymbol{\ell}_0) = (\mathbf{r} \cdot \boldsymbol{\ell}_0)^\eta$ respectively, with $\mathbf{r} = 2(\mathbf{n} \cdot \mathbf{e}) - \mathbf{e}$ the mirror view direction and $\eta \in [0, \infty)$ a shininess parameter.

The main difference between shading terms resides in the choice of reflectance mapping function. Since Phong lobes are defined in the $[0, 1]$ range, the most natural choice is to use them directly as mapping functions: $\delta_j = \rho_j$. It not only identifies a reference direction in which reflected light intensity will be maximal (e.g., \mathbf{n} for δ_d or \mathbf{r} for δ_s), but also provides a natural non-linear fall-off away from this direction. Each term is also enhanced independently with individual scaling magnitudes γ_a, γ_d and γ_s .

Figure 4-a shows results obtained with the scaled Phong Shading model using a single directional light (performances are reported inside the Figure). With such a minimal illumination, the depiction of curvature anisotropy becomes much more sensible; we thus usually make use of low λ values in these settings. Scaling the ambient term gives results equivalent to mean-curvature shading [Kindlmann et al. 2003] (see Figure 4-b). Our method is also easily applied to Toon Shading: one only has to quantize the scaled reflected intensity. However, this quantization tends to mask subtle shading variations, and hence the effectiveness of Radiance Scaling is a bit reduced in this case. Nevertheless, as shown in Figure 4-c, many surface details are still properly enhanced by the technique. We also applied our method to objects made of sub-Lambertian materials ($\rho_{sl}(\boldsymbol{\ell}_0) = (\mathbf{n} \cdot \boldsymbol{\ell}_0)^\zeta, \zeta \in [0, 1]$, with $\delta_{sl} = \rho_{sl}$). Figure 4-d illustrates this process with a sub-Lambertian moon ($\zeta = 0.5$) modeled as a smooth sphere with a detailed normal map.

To test our method in a video game context, we implemented an optimized version of Radiance Scaling using a single light source and Phong shading, and measured an overhead of 0, 17 milliseconds per frame in 1024×768 . Note that our technique is output-sensitive, hence this overhead is independent of scene complexity.

5.2 Complex lighting with Ashikhmin BRDF model

Rendering in complex lighting environments with accurate material models may be done in a variety of ways. In our experiments, we evaluate Ashikhmin’s BRDF model [Ashikhmin et al. 2000] using a dense sampling of directions at each surface point. As for Phong shading, we introduce reflectance mapping functions that let users control the enhancement of different shading terms independently.

Equation 2 has a number of interesting properties, as can be seen in Figure 2. First note that the function is equal to 1 only at $\delta = \alpha$ or when $\kappa = 0$ as required. Second, concave and convex features have a reciprocal effect on the scaling function: $\sigma_{\alpha, \gamma}(\kappa, \delta) = 1/\sigma_{\alpha, \gamma}(-\kappa, \delta)$. A third property is that the function is symmetric with respect to α : $\sigma_{\alpha, \gamma}(\kappa, 1 - \delta) = 1/\sigma_{1-\alpha, \gamma}(\kappa, \delta)$. These choices make the manipulation of the scaling function comprehensible for the user, as illustrated in Figure 3.

Our choice for the curvature mapping function κ is based on the view-centered curvature tensor \mathbf{H} of Vergne et al. [2009]. In the general case, we employ an isotropic curvature mapping: mean curvature is mapped to the $[-1, 1]$ range via $\kappa(\mathbf{p}) = \tanh(\kappa_u + \kappa_v)$ where κ_u and κ_v are the principal curvatures of $\mathbf{H}(\mathbf{p})$. However, for more advanced control, we provide an anisotropic curvature mapping, whereby κ is defined as a function of $\boldsymbol{\ell}$ as well:

$$\kappa(\mathbf{p}, \boldsymbol{\ell}) = \tanh((H + \lambda \Delta_\kappa) \ell_u^2 + (H - \lambda \Delta_\kappa) \ell_v^2 + H \ell_z^2)$$

with the light direction $\boldsymbol{\ell} = (\ell_u, \ell_v, \ell_z)$ expressed in the $(\mathbf{u}, \mathbf{v}, \mathbf{z})$ reference frame, where \mathbf{u} and \mathbf{v} are the principal directions of \mathbf{H} and \mathbf{z} is the direction orthogonal to the picture plane. $H = \kappa_u + \kappa_v$ corresponds to mean curvature and $\Delta_\kappa = \kappa_u - \kappa_v$ is a measure for curvature anisotropy.

Intuitively, the function outputs a curvature value that is obtained by linearly blending principal and mean curvatures based on the projection of $\boldsymbol{\ell}$ in the picture plane. The parameter $\lambda \in [-1, 1]$ controls the way anisotropy is taken into account: when $\lambda = 0$, warping is isotropic ($\forall \boldsymbol{\ell}, \kappa(\boldsymbol{\ell}) = H$); when $\lambda = 1$, warping is anisotropic (e.g., $\kappa(\mathbf{u}) = \kappa_u$); and when $\lambda = -1$, warping is anisotropic, but directions are reversed (e.g., $\kappa(\mathbf{u}) = \kappa_v$). Note however that when $\boldsymbol{\ell}$ is aligned with \mathbf{z} , its projection onto the image plane is undefined, and thus only isotropic warping may be applied ($\forall \lambda, \kappa(\mathbf{z}) = H$).

Radiance Scaling is thus controlled by three parameters: α, γ and λ . The supplementary video illustrates the influence of each parameter on the enhancement effect.

216
217
218
219
220
221
222
223
224
225
226
227

183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215




Figure 4: Radiance Scaling in simple lighting scenarios: (a) Each lobe of Phong’s shading model is scaled independently to reveal shape features such as details in the hair. (b) Radiance Scaling is equivalent to Mean Curvature Shading when applied to an ambient lobe (we combine it with diffuse shading in this Figure). (c) Surface features are also convincingly enhanced with Cartoon Shading, as with this little girl character (e.g., observe the right leg and foot, the ear, the bunches, or the region around the nose). (d) Radiance Scaling is efficient even with sub-Lambertian materials, as in this example of a moon modeled by a sphere and a detailed normal map.

271 Using N light sources and Ashikhmin’s BRDF, Equation 1 becomes

$$L'(\mathbf{p} \rightarrow \mathbf{e}) = \sum_{i=1}^N (\rho_d(\mathbf{l}_i)\sigma_d(\mathbf{p}, \mathbf{l}_i) + \rho_s(\mathbf{e}, \mathbf{l}_i)\sigma_s(\mathbf{p}, \mathbf{e}, \mathbf{l}_i)) L(\mathbf{l}_i)$$

272 where \mathbf{l}_i is the i -th light source direction at point \mathbf{p} and ρ_d and ρ_s
 273 correspond to the diffuse and specular lobes of Ashikhmin’s BRDF
 274 model (see [Ashikhmin et al. 2000]).

275 As opposed to Phong’s model, the diffuse and specular lobes of
 276 Ashikhmin’s BRDF model may be outside of the $[0, 1]$ range, hence
 277 they cannot be used directly as mapping functions. Our alternative
 278 is to rely on each lobe’s reference direction to compute reflectance
 279 mapping functions. We thus choose $\delta_d(\mathbf{l}_i) = (\mathbf{l}_i \cdot \mathbf{n})$ for the diffuse
 280 term and $\delta_s(\mathbf{e}, \mathbf{l}_i) = (\mathbf{h}_i \cdot \mathbf{n})$ for the specular term, where \mathbf{h}_i is the
 281 half vector between \mathbf{l}_i and the view direction \mathbf{e} . As before, each
 282 term is enhanced with separate scaling magnitudes γ_d and γ_s .

283 Figure 5 illustrates the use of Radiance Scaling on a glossy object
 284 with Ashikhmin’s model and an environment map (performances are
 285 reported in Section 6.1). First, the diffuse component is enhanced as
 286 shown in Figure 5-b: observe how concavities are darkened on the

287 chest, the arms, the robe and the hat. The statue’s face gives here
 288 a good illustration of how shading variations are introduced: the
 289 shape of the eyes, mouth and forehead wrinkles is more apparent
 290 because close concavities and convexities give rise to contrasted
 291 diffuse gradients. Second, the specular component is enhanced as
 292 shown in Figure 5-c: this makes the inscriptions on the robe more
 293 apparent, and enhances most of the details on the chest and the
 294 hat. Combining both enhanced components has shown in Figure 5-
 295 d produces a crisp depiction of surface details, while at the same
 296 time conserving the overall object appearance.

5.3 Precomputed radiance data

298 Global illumination techniques are usually time-consuming pro-
 299 cesses. For this reason, various methods have been proposed to
 300 precompute and reuse radiance data. Radiance Scaling introduces
 301 an additional term, σ , to the reflected radiance equation (see Equa-
 302 tion 1). In the general case σ depends both on a curvature mapping
 303 function $\kappa(\mathbf{p})$ and a reflectance mapping function $\delta(\mathbf{e}, \mathbf{l})$, which
 304 means that precomputing *enhanced* radiance data would require at
 305 least an additional storage dimension.

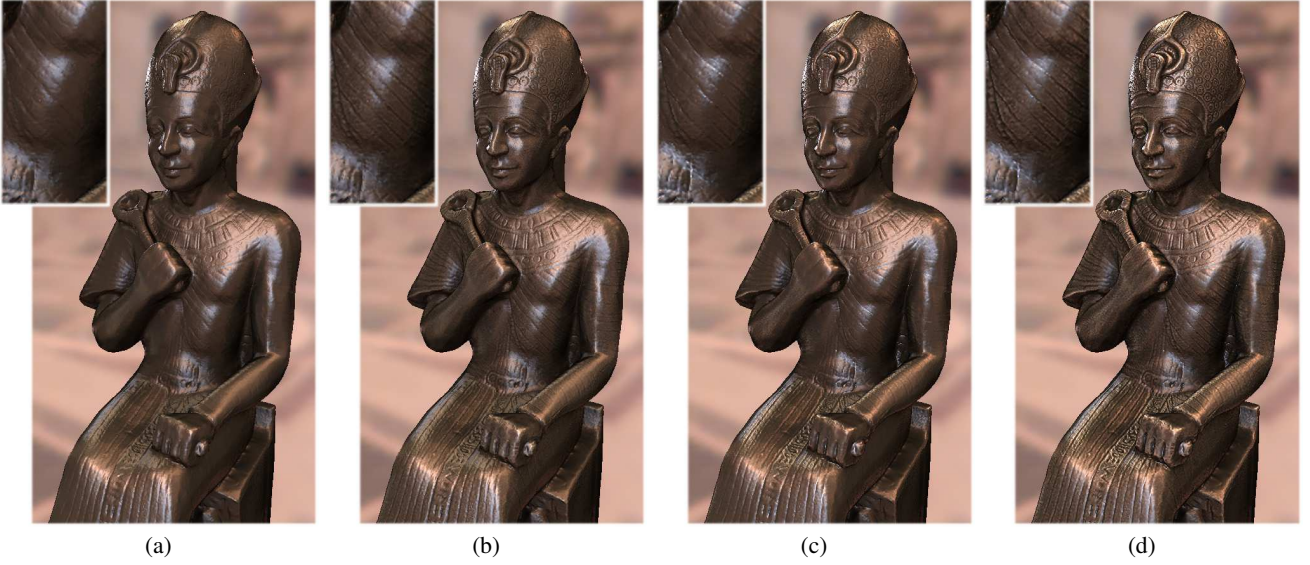


Figure 5: Radiance Scaling using complex lighting: (a) A glossy object obtained with Ashikmin’s BRDF model, with a zoomed view on the chest. (b) Applying Radiance Scaling only to the diffuse term mostly enhances surface features away from highlights (e.g., it darkens concave stripes on the arms and chest). (c) Applying it only to the specular term enhances surface features in a different way (e.g., it brightens some of the concave stripes, and enhances foreshortened areas). (d) Combining both enhancements brings up all surface details in a rich way (e.g., observe the alternations of bright and dark patterns on the chest).

To avoid additional storage, we replace the general reflectance mapping function $\delta(\mathbf{e}, \boldsymbol{\ell})$ by a simplified one $\bar{\delta}(\mathbf{e})$ which is independent of lighting direction $\boldsymbol{\ell}$. The scaling function $\sigma_{\alpha, \gamma}(\kappa(\mathbf{p}), \delta(\mathbf{e}, \boldsymbol{\ell}))$ is then replaced by a simplified version $\bar{\sigma}_{\alpha, \gamma}(\kappa(\mathbf{p}), \bar{\delta}(\mathbf{e}))$, noted $\bar{\sigma}(\mathbf{p}, \mathbf{e})$ and taken out of the integral in Equation 1:

$$L'(\mathbf{p} \rightarrow \mathbf{e}) = \bar{\sigma}(\mathbf{p}, \mathbf{e}) \int_{\Omega} \rho(\mathbf{e}, \boldsymbol{\ell}) (\mathbf{n} \cdot \boldsymbol{\ell}) L(\mathbf{p} \leftarrow \boldsymbol{\ell}) d\boldsymbol{\ell} \quad (3)$$

Now the integral may be precomputed, and the result scaled. Even if scaling is not performed per incoming light direction anymore, it does depend on the curvature mapping function κ , and diffuse and specular components may be manipulated separately by defining dedicated reflectance mapping functions $\bar{\delta}_d$ and $\bar{\delta}_s$. In Sections 5.3.1 and 5.3.2, we show examples of such functions for perfect diffuse, and perfect reflective/refractive materials respectively. The exact same reflectance mapping functions could be used with more complex precomputed radiance transfer methods.

5.3.1 Perfectly diffuse materials

For diffuse materials, *Ambient Occlusion* [Pharr and Green 2004] and *Prefiltered Environment Maps* [Kautz et al. 2000] are among the most widely used techniques to precompute radiance data. We show in the following a similar approximation used in conjunction with Radiance Scaling. The BRDF is first considered constant diffuse: $\rho(\mathbf{e}, \boldsymbol{\ell}) = \rho_d$. We then consider only direct illumination from an environment map: $L(\mathbf{p} \leftarrow \boldsymbol{\ell}) = V(\boldsymbol{\ell}) L_{env}(\boldsymbol{\ell})$ where $V \in \{0, 1\}$ is a visibility term and L_{env} is the environment map. Equation 3 then becomes:

$$L'(\mathbf{p} \rightarrow \mathbf{e}) = \bar{\sigma}(\mathbf{p}, \mathbf{e}) \rho_d \int_{\Omega} (\mathbf{n} \cdot \boldsymbol{\ell}) V(\boldsymbol{\ell}) L_{env}(\boldsymbol{\ell}) d\boldsymbol{\ell}$$

We then approximate the enhanced radiance with

$$L'(\mathbf{p} \rightarrow \mathbf{e}) \simeq \bar{\sigma}(\mathbf{p}, \mathbf{e}) \rho_d A(\mathbf{p}) \bar{L}(\mathbf{n})$$

with $A(\mathbf{p})$ the ambient occlusion stored at each vertex, and \bar{L} an irradiance average stored in a prefiltered environment map:

$$A(\mathbf{p}) = \int_{\Omega} (\mathbf{n} \cdot \boldsymbol{\ell}) V(\boldsymbol{\ell}) d\boldsymbol{\ell}, \quad \bar{L}(\mathbf{n}) = \int_{\Omega} L_{env}(\boldsymbol{\ell}) d\boldsymbol{\ell}$$

For perfectly diffuse materials, we use the reflectance mapping function $\bar{\delta}_d(\mathbf{p}) = \bar{L}(\mathbf{n})/\bar{L}^*$, with \mathbf{n} the normal at \mathbf{p} , and $\bar{L}^* = \max_{\mathbf{n}} \bar{L}(\mathbf{n})$ the maximum averaged radiance found in the prefiltered environment map. This choice is coherent with perfectly diffuse materials, since in this case the light direction that contributes the most to reflected light intensity is the normal direction on average.

Figure 6-a shows the warping of prefiltered environment maps using the Armadillo model. Observe how macro-geometry patterns are enhanced on the leg, arm and forehead. The ambient occlusion term is shown separately in Figure 6-b. An alternative to using a prefiltered environment map for stylized rendering purpose is the Lit Sphere [Sloan et al. 2001]. It consists in a painted sphere where material, style and illumination direction are implicitly given, and has been used for volumetric rendering [Bruckner and Gröller 2007] and in the ZBrush® software (under the name “matcap”). Radiance Scaling produces convincing results with Lit Spheres as shown in Figure 6-c and in the supplementary video.

5.3.2 Perfectly reflective and refractive materials

The case of perfectly reflective or refractive materials is quite similar to the perfectly diffuse one. If we consider a perfectly reflective/refractive material ρ_s (a dirac in the reflected/refracted direction \mathbf{r}) and ignore the visibility term, then Equation 3 becomes:

$$L'(\mathbf{p} \rightarrow \mathbf{e}) = \bar{\sigma}(\mathbf{p}, \mathbf{e}) L_{env}(\mathbf{r})$$

We use the reflectance mapping function $\bar{\delta}_s(\mathbf{e}) = L_{env}(\mathbf{r})/L_{env}^*$, with \mathbf{r} the reflected/refracted view direction and $L_{env}^* = \max_{\mathbf{r}} L_{env}(\mathbf{r})$ the maximum irradiance in the environment map. This choice is coherent with perfectly reflective/refractive materials, since in this case the light direction that contributes the most to reflected light intensity is the reflected/refracted view direction.




Figure 6: Radiance Scaling using precomputed lighting: (a) To improve run-time performance, precomputed radiance data may be stored in the form of ambient occlusion and prefiltered environment map. Radiance Scaling is easily adapted to such settings and provides enhancement at real-time frames rates (66 fps /345,944 polygons). (b) Even when only applied to the ambient occlusion term, Radiance Scaling produces convincing results. (c) For stylized rendering purposes, Radiance Scaling may be applied to a Lit Sphere rendering.

362 Figure 1-right shows how Radiance Scaling enhances surface features with a simple approximation of a purely refractive material.
 363 The video also shows results when the method is applied to an object with a mirror-like material.
 364
 365

366 6 Discussion

367 We first compare Radiance Scaling with previous methods in Section 6.1, with a focus on Light Warping [Vergne et al. 2009] since it relies on the same surface features. We then discuss limitations and avenues for future work in Section 6.2.
 368
 369
 370

371 6.1 Comparisons with previous work

372 Our approach is designed to depict local surface features, and is difficult to compare with approaches such as Accessibility Shading that consider more of the surrounding geometry. Accessibility Shading characterizes how easily a surface may be touched by a spherical probe, and thus tends to depict more volumetric features. However, for surfaces where small-scale relief dominates large-scale variations (such as carved stones or roughly textured statues), the spherical probe acts as a curvature measure. In this case, Accessibility Shading becomes similar to Mean Curvature Shading, which is a special case of Radiance Scaling as seen in Figure 4-b.
 373
 374
 375
 376
 377
 378
 379
 380
 381

382 A technique related to Accessibility Shading is Ambient Occlusion: indeed, measuring occlusion from visible geometry around a surface point is another way of probing a surface. Ambient Occlusion is more efficient at depicting proximity relations between objects (such as contacts), and deep cavities. However, as seen in Figure 6-b, it also misses shallow (yet salient) surface details, or even smooth them out. Radiance Scaling reintroduces these details seamlessly. Both methods are thus naturally combined to depict different aspects of object shape.
 383
 384
 385
 386
 387
 388
 389
 390

391 3D Unsharp Masking provides yet another mean to enhance shape features: by enhancing outgoing radiance with a Cornsweet illusion effect, object shape properties correlated to shading are enhanced along the way. Besides the fact that users have little control on what property of a scene will be enhanced, 3D Unsharp Masking tends
 392
 393
 394
 395

396 to make flat surfaces appear rounded, as in Cignoni et al. [2005]. It is also limited regarding material appearance, as pointed out in Vergne et al. [2009]. We thus focus on a comparison with Light Warping in the remainder of this Section.
 397
 398
 399

400 An important advantage of Radiance Scaling over Light Warping is that it does not require a dense sampling of the environment illumination, and thus works in simple rendering settings as described in Section 5.1. As an example, consider Toon Shading. Light Warping does allow to create enhanced cartoon renderings, but for this purpose makes use of a minimal environment illumination, and still requires to shoot multiple light rays. Radiance Scaling avoids such unnecessary sampling of the environment as it works with a single light source. Hence it is much faster to render: the character in Figure 4 is rendered at 241 fps with Radiance Scaling, whereas performances drop to 90 fps with Light Warping as it requires at least 16 illumination samples to give a convincing result.
 401
 402
 403
 404
 405
 406
 407
 408
 409
 410
 411

412 For more complex materials, Radiance Scaling is also faster than




Figure 7: This plot gives the performances obtained with the scene shown in Figure 5 without enhancement, and with both Radiance Scaling and Light Warping. The 3D model is composed of 1,652,528 polygons. While the time for rendering a single frame increases linearly with the number of light samples in all cases, our novel method is linearly faster than Light Warping.




Figure 8: Comparison with Light Warping. Top row: Light warping image. Bottom row: Radiance Scaling. (a-b) Both methods show similar enhancement abilities when used with a diffuse material and a natural illumination environment: convexities exhibit brighter colors, and concavities darker colors in most cases. For some orientation of the viewpoint relative to the environment, Light Warping may reverse this effect though (concavities are brighter, convexities darker) while Radiance Scaling does not. (c-d) The methods are most different with shiny objects, shown with two illumination orientations as well.

413 Light Warping, as seen in Figure 7. However, the two methods
 414 are not qualitatively equivalent, as shown in Figure 8. For dif-
 415 fuse materials and with natural illumination, the two methods pro-
 416 duce similar results: concavities are depicted with darker colors,
 417 and convexities with brighter colors. However, for some orienta-
 418 tions of the viewpoint relative to the environment illumination,
 419 Light Warping may reverse this effect, since rays are attracted to-
 420 ward or away from the camera regardless of light source locations.
 421 Radiance Scaling does not reverse tone in this manner. The main
 422 difference between the two techniques appears with shiny materi-
 423 als. In this case, the effect of enhancement on illumination is more
 424 clearly visible: Light Warping modulates lighting frequency, while
 425 Radiance Scaling modulates lighting intensity, as is best seen in the
 426 supplementary video.

427 6.2 Directions for future work

428 We have shown that the adjustment of reflected light intensities,
 429 a process we call Radiance Scaling, provides a versatile approach
 430 to the enhancement of surface shape through shading. However,

431 when the enhancement magnitude is pushed to extreme values, our
 432 method alters material appearance. This is because variations in
 433 shape tend to dominate variations due to shading. An exciting av-
 434 enue of future work would be to characterize perceptual cues to
 435 material appearance and preserve them through enhancement.

436 Although Radiance Scaling produces convincing enhancement in
 437 many rendering scenarios, there is still room for alternative en-
 438 hancement techniques. Indeed, our approach makes two assump-
 439 tions that could be dropped in future work: 1) concave and convex
 440 features have inverse effects on scaling; and 2) enhancement is ob-
 441 tained by local differential operators. The class of reflected lighting
 442 patterns humans are able to make use for perceiving shape is obvi-
 443 ously much more diverse than simple alternations of bright and dark
 444 colors in convexities and concavities [Koenderink J.J. 2003]. And
 445 these patterns are likely to be dependent on the main illumination
 446 direction(e.g., [Ho et al. 2006; Caniard and Fleming 2007; O’Shea
 447 et al. 2008]), material characteristics (e.g., [Adelson 2001; Vangorp
 448 et al. 2007]), motion (e.g., [Pont S.C. 2003; Adato et al. 2007]), and
 449 silhouette shape (e.g., [Fleming et al. 2004]). Characterizing such

450 patterns is a challenging avenue of future work.

451 References

- 452 ADATO, Y., VASILYEV, Y., BEN SHAHAR, O., AND ZICKLER, T.
453 2007. Toward a theory of shape from specular flow. In *ICCV07*,
454 1–8.
- 455 ADELSON, E. H. 2001. On seeing stuff: the perception of materi-
456 als by humans and machines. In *Society of Photo-Optical Instru-*
457 *mentation Engineers (SPIE) Conference Series*, B. E. Rogowitz
458 and T. N. Pappas, Eds., vol. 4299 of *Presented at the Society*
459 *of Photo-Optical Instrumentation Engineers (SPIE) Conference*,
460 1–12.
- 461 ASHIKHMIN, M., PREMOZE, S., AND SHIRLEY, P. 2000. A
462 microfacet-based BRDF generator. In *Proc. ACM SIGGRAPH*
463 *'00*, ACM, 65–74.
- 464 BRUCKNER, S., AND GRÖLLER, M. E. 2007. Style transfer func-
465 tions for illustrative volume rendering. *Computer Graphics For-*
466 *um* 26, 3 (Sept.), 715–724.
- 467 CANIARD, F., AND FLEMING, R. W. 2007. Distortion in 3D shape
468 estimation with changes in illumination. In *APGV '07: Proc.*
469 *symposium on Applied perception in graphics and visualization*,
470 ACM, 99–105.
- 471 CIGNONI, P., SCOPIGNO, R., AND TARINI, M. 2005. A
472 simple Normal Enhancement technique for Interactive Non-
473 photorealistic Renderings. *Comp. & Graph.* 29, 1, 125–133.
- 474 DECARLO, D., FINKELSTEIN, A., RUSINKIEWICZ, S., AND
475 SANTELLA, A. 2003. Suggestive Contours for Conveying
476 Shape. *ACM Trans. Graph. (Proc. SIGGRAPH 2003)* 22, 3, 848–
477 855.
- 478 FLEMING, R. W., TORRALBA, A., AND ADELSON, E. H. 2004.
479 Specular reflections and the perception of shape. *J. Vis.* 4, 9 (9),
480 798–820.
- 481 GOODWIN, T., VOLLICK, I., AND HERTZMANN, A. 2007.
482 Isophote distance: a shading approach to artistic stroke thick-
483 ness. In *NPAR '07: Proc. international symposium on Non-*
484 *photorealistic animation and rendering*, ACM, 53–62.
- 485 HO, Y.-X., LANDY, M. S., AND MALONEY, L. T. 2006. How di-
486 rection of illumination affects visually perceived surface rough-
487 ness. *J. Vis.* 6, 5 (5), 634–648.
- 488 JUDD, T., DURAND, F., AND ADELSON, E. H. 2007. Appar-
489 ent Ridges for Line Drawing. *ACM Trans. Graph. (Proc. SIG-*
490 *GRAPH 2007)* 26, 3, 19.
- 491 KAUTZ, J., VÁZQUEZ, P.-P., HEIDRICH, W., AND SEIDEL, H.-
492 P. 2000. Unified Approach to Prefiltered Environment Maps. In
493 *Proceedings of the Eurographics Workshop on Rendering Tech-*
494 *niques 2000*, Springer-Verlag, 185–196.
- 495 KINDLMANN, G., WHITAKER, R., TASHIZEN, T., AND MÖLLER,
496 T. 2003. Curvature-Based Transfer Functions for Direct Volume
497 Rendering: Methods and Applications. In *Proc. IEEE Visualiza-*
498 *tion 2003*, 513–520.
- 499 KOENDERINK J.J., D. A. v. 2003. *The visual neurosciences*. MIT
500 Press, Cambridge, ch. Shape and shading, 1090–1105.
- 501 KOLOMENKIN, M., SHIMSHONI, I., AND TAL, A. 2008. Demar-
502 cating Curves for Shape Illustration. *ACM Trans. Graph. (Proc.*
503 *SIGGRAPH Asia 2008)* 27, 5, 1–9.
- 504 LEE, Y., MARKOSIAN, L., LEE, S., AND HUGHES, J. F. 2007.
505 Line drawings via abstracted shading. *ACM Trans. Graph.* 26, 3,
506 18.
- 507 MILLER, G. 1994. Efficient Algorithms for Local and Global
508 Accessibility Shading . In *Proc. ACM SIGGRAPH '94*, ACM,
509 319–326.
- 510 NIENHAUS, M., AND DÖLLNER, J. 2004. Blueprints: illustrating
511 architecture and technical parts using hardware-accelerated non-
512 photorealistic rendering. In *Graphics Interface (GI'04)*, Cana-
513 dian Human-Computer Communications Society, 49–56.
- 514 OHTAKE, Y., BELYAEV, A., AND SEIDEL, H.-P. 2004. Ridge-
515 valley lines on meshes via implicit surface fitting. *ACM Trans.*
516 *Graph. (Proc. SIGGRAPH 2004)* 3, 23, 609–612.
- 517 O'SHEA, J. P., BANKS, M. S., AND AGRAWALA, M. 2008. The
518 assumed light direction for perceiving shape from shading. In
519 *APGV '08: Proc. symposium on Applied perception in graphics*
520 *and visualization*, ACM, 135–142.
- 521 PHARR, M., AND GREEN, S. 2004. *GPU Gems*. Addison-Wesley,
522 ch. Ambient Occlusion.
- 523 PONT S.C., K. J. 2003. *Computer Analysis of Images and Patterns*.
524 Springer, Berlin, ch. Illuminance flow, 90–97.
- 525 RITSCHHEL, T., SMITH, K., IHRKE, M., GROSCH, T.,
526 MYSZKOWSKI, K., AND SEIDEL, H.-P. 2008. 3D Unsharp
527 Masking for Scene Coherent Enhancement. *ACM Trans. Graph.*
528 *(Proc. SIGGRAPH 2008)* 27, 3, 1–8.
- 529 RUSINKIEWICZ, S., BURNS, M., AND DECARLO, D. 2006. Ex-
530 aggerated Shading for Depicting Shape and Detail. *ACM Trans.*
531 *Graph. (Proc. SIGGRAPH 2006)* 25, 3, 1199–1205.
- 532 SAITO, T., AND TAKAHASHI, T. 1990. Comprehensible Ren-
533 dering of 3-D Shapes. In *Proc. ACM SIGGRAPH '90*, ACM,
534 197–206.
- 535 SLOAN, P.-P. J., MARTIN, W., GOOCH, A., AND GOOCH, B.
536 2001. The lit sphere: A model for capturing NPR shading from
537 art. In *Graphics interface 2001*, Canadian Information Process-
538 ing Society, 143–150.
- 539 VANGORP, P., LAURIJSSSEN, J., AND DUTRÉ, P. 2007. The in-
540 fluence of shape on the perception of material reflectance. *ACM*
541 *Trans. Graph. (Proc. SIGGRAPH 2007)* 26, 3, 77.
- 542 VERGNE, R., BARLA, P., GRANIER, X., AND SCHLICK, C. 2008.
543 Apparent relief: a shape descriptor for stylized shading. In *NPAR*
544 *'08: Proc. international symposium on Non-photorealistic ani-*
545 *mation and rendering*, ACM, 23–29.
- 546 VERGNE, R., PACANOWSKI, R., BARLA, P., GRANIER, X., AND
547 SCHLICK, C. 2009. Light warping for enhanced surface depic-
548 tion. *ACM Transaction on Graphics (Proceedings of SIGGRAPH*
549 *2009)* (Aug).
- 550 ZHANG, L., HE, Y., XIE, X., AND CHEN, W. 2009. Laplacian
551 Lines for Real Time Shape Illustration. In *I3D '09: Proc. sym-*
552 *posium on Interactive 3D graphics and games*, ACM.