

Local Light Alignment for Multi-Scale Shape Depiction

Nolan Mestres¹ Romain Vergne¹ Camille Noûs² Joëlle Thollot¹

¹ Univ. Grenoble Alpes, CNRS, Inria, Grenoble INP, LJK
² Cogitamus Laboratory



Figure 1: Multi-scale shape depiction enhancements using our approach in real-time (left pair) and global illumination (right pair) scenarios. The shield is rendered using Cook-Torrance model [CT82]. Note how shape features of various sizes are enhanced without impairing lighting and material appearance.

Abstract

Motivated by recent findings in the field of visual perception, we present a novel approach for enhancing shape depiction and perception of surface details. We propose a shading-based technique that relies on locally adjusting the direction of light to account for the different components of materials. Our approach ensures congruence between shape and shading flows, leading to an effective enhancement of the perception of shape and details while impairing neither the lighting nor the appearance of materials. It is formulated in a general way allowing its use for multiple scales enhancement in real-time on the GPU, as well as in global illumination contexts. We also provide artists with fine control over the enhancement at each scale.

CCS Concepts

• Computing methodologies → Non-photorealistic rendering; Perception;

1. Introduction

Producing an image able to efficiently convey the shape of objects in a scene is crucial in many domains such as medical imaging, cartography, archaeology, video games or film industry. Our work targets designers exploring or creating visualizations from 3D scenes, as they have a strong need for techniques enhancing the depiction of shape. The computer graphics community has tackled this challenge with increasing success in the past thirty years but it still is an open field of research. Inspired by recent studies and analyses related to the perception of shape-from-shading, our goal in this paper is to better depict shape features at different scales without impairing the final look of the scene.

Since the seminal work of Koenderink and Van Doorn [KvD80] on isophote patterns, more and more evidence shows that the Hu-

man Visual System (HVS) is using the so-called *shading flows* (i.e. directions and magnitudes of shading patterns in images) to estimate 3D shape information [BSZ01, HCBZ01, FTA04]. These patterns are correlated to view-dependent second-order shape properties and therefore provide information about the curvature of the surface, regardless of the material and illumination conditions [VBBF16]. The relationships between *shading flows* and *shape flows* have proven to be effective in the vision [KHRA*18] as well as in the graphics [VBF12] communities to estimate or control perceived shape of objects.

However, this correlation is not perfect. Depending on lighting conditions or material properties, shading flows may deviate from the physical shape and then lead to misperceived surface features. For instance, strong directional lighting tends to mask and distort

the perception of shape [CF07]. The perceived surface of objects can also be attenuated or exaggerated depending on their materials [WDKP12, MA14]. Marlow *et al.* [MKA11] also studied how a lack of congruence between diffuse and specular flows has a negative impact on our perception of gloss.

Motivated by these findings in visual perception, we propose a simple yet effective method that locally adjusts light directions in such a way that the congruence between shape and shading flows is increased, ensuring a better depiction of shape features in rendered images (Figure 1). Our key idea is to align the light with shape features at multiple scales to better control the final contrast and direction of the shading patterns. With the HVS being not very sensitive to lighting inconsistencies [OCS01, LSSG10], we show how to properly modify directions of light to obtain spatially and temporally coherent depicted shapes (Sections 4, 5). Our approach can be integrated into any rendering algorithm (*i.e.* from real-time to offline) and easily adapts to any component of the shading function or the BRDF, to control the behavior of highlights for instance (Section 6). Through its fine multi-scale controls, our method provides end-users with efficient means to depict arbitrary features of an input shape, independently of materials and lighting conditions.

2. Related work

For over three decades the computer graphics community proposed many methods to enhance the depiction of 3D shapes. Two main classes emerged: line-based rendering [ST90, DFRS03, JDA07, KST08] which focuses on finding the proper set of lines to depict shape, and shading-based techniques that aim at modifying the shading in order to better depict shape. Our approach falls into the latter. Even though line-based techniques make no use of shading information (exception made of Lee *et al.* [LMLH07]), they could still complement our approach.

Contrast enhancement. Image contrast enhancement methods pursue in essence the same goal as we do: enhancing the contrast to better depict the shape of the subject in the final image [Sta00, AKAC07, RRAAW*16]. Most of these image processing techniques are based on histogram equalization and use 2D information from an input photograph. Even though they are not real-time, they could be used as a post-process operation on rendered images as well, and still achieve convincing results. However, we aim for real-time multi-scale control over the shape depiction of objects in 3D scenes.

Accessibility Shading methods [Mil94, PG04] were the first to locally adapt shading. They focus on emphasizing concavities based on information about the surrounding geometry. A well known example is Ambient Occlusion, still widely used for real-time rendering as an approximation for global illumination.

Cignoni *et al.* [CST05] introduced a technique called Normal Enhancement that relies on adding a low-passed filtered version of the normal field to the original normals of the object. It enhances the high frequencies in the normal field, thus effectively increasing the perception of contours. This technique requires a preprocess on mesh vertices, making its use with dynamic scenes difficult. It also tends to impair the perception of planar surfaces, exhibiting concavity artifacts.

Later, Ritschel *et al.* [RSI*08] used the same concept for their 3D Unsharp Masking method but applied to the outgoing radiance in the scene. It allows the use of arbitrary illumination and materials, even if some halo artifacts can occur with glossy surfaces. By unsharp masking the outgoing radiance, they effectively increase the contrast in the scene among other cues. The shape enhancement is thus implicit and occurs when surface shape and radiance are correlated.

Radiance Scaling [VPB*11] introduces a novel function to control the incident radiance in the rendering equation. By enhancing shape depiction in a way that is both dependent on surface curvature and reflectance this method addresses most of the drawbacks of the previous work making the enhancement explicit. However, this approach only works at a single scale and may produce under/over-saturated results that alter material appearance.

Moving light(s). An alternative to directly modifying the output radiance in a scene is to locally adjust the direction of lights to make sure that shapes in rendered images are properly depicted.

Light Warping [VPB*09] aims at enhancing concavities and convexities by stretching or compressing the light environment based on a view-dependent analysis of the surface curvature. The technique depends on the statistics of the scene, that must exhibit enough variations in contrast. As directions of light sources are not taken into account by the warping operation, it has sometimes the side effect of inverting bright and dark regions.

Exaggerated Shading [RBD06] locally adjusts the light so that it is always at grazing angles with respect to the surface. This method is effective at depicting shape features but suffers from two main limitations: adjusting only the elevation of the light source prevents some features to be exaggerated (*e.g.* ridges aligned with the light); small features may not produce enough contrast to be well depicted, even if the light is located at grazing angles. To get around these problems and ensure strong contrast everywhere, the authors rely on a modified half-Lambertian model, making the technique tied to a unique shading function. Unlike previous approaches, Exaggerated Shading accounts for multiple scales enhancements: renderings are computed for each detail scale, then weighted and added together to obtain the final image, allowing for a fine control over the result.

Taking inspiration from Exaggerated Shading for the multi-scale control, but targeting arbitrary materials as with Light Warping, we propose to locally adjust the direction of light so as to increase the congruence between shape and shading flows.

3. Approach overview

Our main goal is to enhance shape details at multiple scales by controlling the intensity of the shading regardless of the shading function. Intuitively we propose to locally rotate the light to ensure that, for each surface feature, the intensity is maximal on one side, and minimal on the other (see Figure 2).

By feature we mean any significant part of a shape, typically a ridge or a valley. Such features are scale-dependent: a very small ridge is a significant feature at a fine scale but should be omitted at

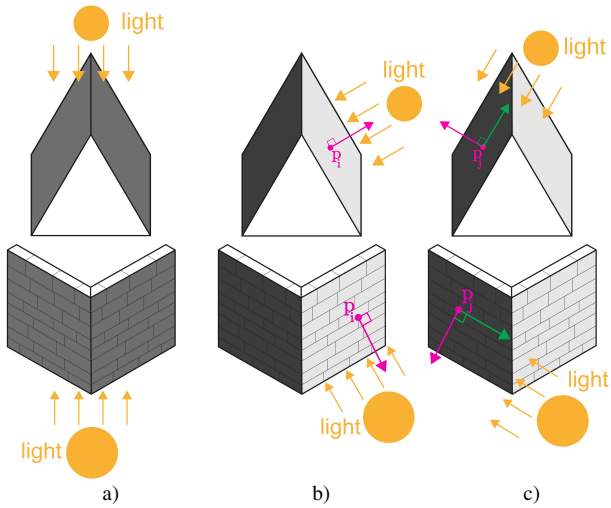


Figure 2: 3D isometric view of a ridge (top row) and a corner wall (bottom row). When the ridge is lit from above, or the wall frontally, by a strong directional light, both sides have the same shading (a), thus making the shape difficult to understand. For the shading to exhibit contrast, we change the vertical and horizontal directions of light at each point \mathbf{p} of the surface, so that the light faces the normal vector (pink) on one side of the feature (b) and is perpendicular to the normal (green) on the other (c). It ensures having bright and dark sides, resulting in a better perceived shape.

a coarser one. Therefore when applied at multiple scales our light adjustment allows us to enhance the perception of the overall shape while still revealing even the smallest details of the model.

Our approach works in two steps:

- Creating local frames at multiple scales: this is done by successively smoothing the normal field of the object. Surface features are to be understood as the geometry of the model that varies between two successive versions of the normals.
- Adjusting the light direction: we design a function that locally adjusts the incident direction of light with respect to the geometry of features at each point. We perform this step iteratively using two successive levels of the normals.

In the following we detail how to modify the direction of light in the case of a single scale (Section 4) before showing how to make this light adjustment controllable at multiple scales (Section 5). We then generalize the enhancement process to account for arbitrary materials (Section 6) and show its use with various rendering techniques from real-time rendering to global illumination (Section 7, Results). Finally we analyze the congruence between shape and shading provided by our approach and show comparisons with previous techniques (Section 7, Congruence score, Comparison with previous work).

4. A Lambertian feature on a plane

Let's consider a ridge lit by a single directional light (Figure 2). Our idea is somewhat similar in essence to that of Exaggerated Shading [RBD06], with a major difference: we do not seek to place the




Figure 3: Local space $\mathcal{L} = (\mathbf{x}, \mathbf{y}, \mathbf{z})$ used for the computation of θ at point \mathbf{p} . The up vector \mathbf{z} defines the orientation of the surface at a coarser scale. \mathbf{x} is the (renormalized) projection of the normal \mathbf{n} on the base plane, and $\mathbf{y} = \mathbf{x} \times \mathbf{z}$. $\mathcal{R}_{(\mathbf{a}, \theta)}$ denotes the resulting rotation.

light at a grazing angle with respect to the shape detail, but rather to maximize the difference in intensity between the shading of both of its sides. Considering simple Lambertian shading, it means aligning the light with the normal on one side, and with a vector orthogonal to the normal on the other one (i.e. a properly chosen tangent vector) so that the scalar products between lights and normals are always equal to either 1 or 0 on a feature. What we call a feature is an element of interest of the shape protruding from a coarser version of the surface, later referenced as the base. The base depends on the scale-space analysis and will be detailed in the next section.

Light rotation. At a given point \mathbf{p} of the surface, let's consider a local reference frame $\mathcal{L} = (\mathbf{x}, \mathbf{y}, \mathbf{z})$, where \mathbf{z} is the base plane up vector, \mathbf{x} is the projection of the normal \mathbf{n} onto the base plane and $\mathbf{y} = \mathbf{x} \times \mathbf{z}$, as shown in Figure 3. Assuming that both the light direction \mathbf{l} and the normal vector \mathbf{n} are expressed in the local frame \mathcal{L} , our goal is to define a guiding vector \mathbf{g} , towards which the light will be directed to increase the contrast of the underlying feature. In a diffuse scenario we define \mathbf{g} as either the normal for the bright side of the feature or the tangent that lies in the (\mathbf{x}, \mathbf{z}) plane for the dark side:

$$\mathbf{g} = \begin{cases} \mathbf{n}, & \text{if } (\bar{\mathbf{l}} \cdot \mathbf{x}) \geq 0 \\ \mathcal{R}_{(\mathbf{y}, -\frac{\pi}{2})} \mathbf{n}, & \text{otherwise,} \end{cases}$$

where $\mathcal{R}_{(\mathbf{y}, -\frac{\pi}{2})}$ denotes a rotation by $-\frac{\pi}{2}$ around the \mathbf{y} axis and $\bar{\mathbf{v}}$ represents the projection of a vector \mathbf{v} onto the (\mathbf{x}, \mathbf{y}) plane. The original direction of light is then rotated towards \mathbf{g} :

$$\begin{aligned} \mathbf{l}' &= \mathcal{R}_{(\mathbf{a}, \theta)} \mathbf{l}, & \text{with} \\ \mathbf{a} &= \mathbf{l} \times \mathbf{g}, \\ \theta &= \sigma \cdot W_{\gamma}(\lambda) \cdot \hat{\mathbf{l}} \cdot \mathbf{g}, \end{aligned} \tag{1}$$

with $\sigma \in [0, 1]$ a user-defined parameter that controls the strength of the enhancement. Aligning the light with \mathbf{g} leads to binary images as shown in Figure 4, b. $W_{\gamma} : [0, 1] \rightarrow [0, 1]$ is therefore a weighting function that prevents discontinuities to appear when the light is switching from bright to dark regions. W_{γ} is function of a confidence value λ that represents how close we are from a disconti-




Figure 4: Effects of λ_1 and λ_2 . a) The light is not adjusted; b) when $\lambda = 1$ the adjustment is unconstrained, hence the discontinuities; c) λ_1 accounts for the cases where \mathbf{g} changes from the normals to the tangents; d) λ_2 accounts for cases where details are not protruding enough from the base plane; e) using both constraints solves the discontinuities. Images rendered using $\gamma = 2$ and $\sigma = 1$.




Figure 5: Renderings using different γ values, with $\lambda = \lambda_1 \cdot \lambda_2$ and $\sigma = 1$. γ is used to smooth or sharpen the transitions in shading intensity. A high value of gamma results in quick transitions between bright and dark sides of the features, while a small value produces wider transition regions.

nity. There are two scenarios in which the light direction changes suddenly, causing the discontinuities:

- When $(\bar{\mathbf{I}} \cdot \mathbf{x}) = 0$, corresponding to the locations around which the guiding vector \mathbf{g} is switching from the normal to the tangent vectors. We represent it as $\lambda_1 = 1 - 2\widehat{\mathbf{I}\mathbf{g}}/\pi$, the distance from the light to this switching point.
- When there are little to no features to enhance. As the feature normal \mathbf{n} tends to be aligned with the base plane up vector \mathbf{z} , enhancing the feature makes less sense, the local frame \mathcal{L} being almost undefined. We detect this case using the magnitude of the projected normal on the (\mathbf{x}, \mathbf{y}) plane, as $\lambda_2 = \|\bar{\mathbf{n}}\|$.

We finally combine these confidence values in the weighting function W_γ to ensure smooth transitions in the shading everywhere on the surface:

$$W_\gamma(\lambda) = \frac{\lambda}{e^{-\gamma(1-\lambda)} + \lambda}, \quad \text{with } \lambda = \lambda_1 \cdot \lambda_2 \text{ and } \lambda_1, \lambda_2 \in [0, 1].$$

Intuitively, this function controls how fast the light can rotate around discontinuities. When approaching a region where the rotation of \mathbf{I} is ambiguous, λ tends to 0, thus preventing the rotation. Away from the discontinuities, λ quickly reaches 1. Figure 4 illustrates the effect over the discontinuities of both λ_1 and λ_2 when applied separately (c,d) or combined (e). $W_\gamma(\lambda)$ is a simple remapping of the λ value, controllable via the γ parameter. When $\gamma = 0$, the function is linear ($W_0(\lambda) = \lambda$). Increasing γ narrows the transition areas between bright and dark sides, while still ensuring them to be smooth, as seen in Figure 5. In practice and unless mentioned otherwise, $\gamma = 3$ in all the results presented in the paper, as it provides a good trade-off between the resulting contrast and perceived discontinuities.




Figure 6: Illustration of our multi-scale enhancement using $S = 3$ scales, applied on two points \mathbf{p} and \mathbf{q} of a surface. Normals \mathbf{n}_3 of the last scale are used to deduce the local frame \mathcal{L}_2 , only represented by the z -axis (in blue) for clarity: $\mathbf{z}_2 = \mathbf{n}_3$, and similarly, $\mathbf{z}_1 = \mathbf{n}_2$. The direction of light is then rotated towards the guiding vectors at each successive scale. \mathbf{l}_3 becomes \mathbf{l}_2 which in turn becomes \mathbf{l}_1 . The choice of guiding vectors depends on the sign of the scalar product between \mathbf{l}_{i+1} and \mathbf{n}_i when projected onto the plane defined by \mathbf{z}_i (Section 4). This is easily seen here by testing if these vectors are on the same side with respect to the z -axis. At scale 2 for instance, \mathbf{l}_3 and \mathbf{n}_2 are on the same side at point \mathbf{q} or on different sides at point \mathbf{p} with respect to \mathbf{z}_2 , making the surface appear brighter on the right and darker on the left. The same approach is applied at scale 1 to reveal the finest features.

5. Multiple scales enhancement

As opposed to Exaggerated Shading [RBD06], we do not seek to average multiple renderings, but rather to compute a single direction of light that takes all scales into account at once. With S being the number of scales, our rotation operator (Equation 1) is applied successively $S - 1$ times, from the coarsest to the finest scale, to control the depiction of each shape feature:

$$\begin{aligned} \mathbf{l}_S &= \mathbf{l}, \\ \mathbf{l}_i &= \mathcal{R}_{(\mathbf{a}_i, \theta_i)} \mathbf{l}_{i+1}, \quad \text{for } i \in [1, S-1]. \end{aligned}$$

W_γ used in the computation of θ ensures that the light direction is rotated only for points lying on features of the considered scale. For other points, light directions rotated at previous (coarser) scales remain unchanged. If a point belongs to features of different sizes then the light is successively rotated to reveal small details while still retaining the overall direction induced by bigger features.

We iteratively filter the input shape geometry to obtain a set of S normals at each point of the surface. These are denoted \mathbf{n}_i ($i \in [1, S]$). \mathbf{n}_1 and \mathbf{n}_S correspond respectively to the finest and coarsest




Figure 7: Enhancements at arbitrary detail scales (all scales are displayed in the supplementary document). Each row but the last is a detail scale, from the coarsest (5) to the finest (1). Each column shows a different value of σ . The coarsest scales enhance the overall shape while finer scales enhance smaller details of the model. Bottom: model without enhancement (left) and a fine-tuned enhanced version where $S = 6$ with $\sigma_i = \{0.5, 0.3, 0.1, 0.2, 0.2\}$.

versions of the surface. The local reference frame $\mathcal{L}_i = (\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i)$ ($i \in [1, S - 1]$) in which we rotate the light at scale i differs at each scale. It is a convenient way to represent geometric features, *i.e.* variations between two successive smoother versions of the surface normal field. We compute it as:

$$\mathcal{L}_i : \begin{cases} \mathbf{z}_i &= \mathbf{n}_{i+1}, \\ \mathbf{x}_i &= \mathbf{n}_i - \mathbf{z}_i(\mathbf{n}_i \cdot \mathbf{z}_i), \\ \mathbf{y}_i &= \mathbf{x}_i \times \mathbf{z}_i, \end{cases}$$

where each vector is normalized to ensure a proper change of basis to compute the guide and rotate the light direction. An illustration of our multi-scale enhancement is shown in Figure 6 where we can observe how the original light \mathbf{l}_3 is modified at each scale to reveal coarse (scale 2) and fine (scale 1) details.

Implementation details. In practice, we store the input normals in a G-Buffer and use successive bilateral filtering operations [BCCS12] to obtain the set of smoothed normals. The range kernel is applied on depth values to preserve the silhouettes of the surface. The size of the spatial kernels depends on the image resolution and the desired level of detail scales. In our results, we in-




Figure 8: Specular (a, b) and transmitted (c, d) multi-scale enhancements using respectively the reflected and transmitted view vectors as guides, with $S = 5$. The left pair of images displays the specular (resp. transmitted) term only; on the right pair we reintroduced the diffuse term (resp. specular term and background). Our method modifies the shape of highlights and intensity of refractions in such a way that details are revealed at the corresponding scale. Parameters used are $\sigma_i = \{0.4, 0.2, 0, 0.3\}$ (a), $\sigma_i = \{0.4, 0.2, 0.2, 0.3\}$ (b), $\sigma_i = \{0.7, 0.3, 0.25, 0.6\}$ (c, d).

creased this size by a factor of 1.7 for each coarser scale. All the computations are then done per-pixel, in a fragment shader. The resulting modified directions of light are also stored in a buffer, that can next be used as input to any shading function.

We found that a value of $S \in \{5, 6, 7\}$ is mostly enough to enhance the depiction of an arbitrary surface, but of course it is also a matter of artistic choice. Users can control the intensity of the enhancement at each scale as presented in Figure 7. The x-axis shows how σ influences the enhancement while the y-axis corresponds to different scales of the model. When successively applied (bottom-right) our method allows the surface to be properly depicted everywhere. Note that σ may also be tuned differently at each scale to precisely control the rendered image. A designer may then easily and interactively tweak each scale enhancement factor σ_i to explore and create his own intended visual representation.

6. Arbitrary materials enhancement

Our approach can easily be extended to control the properties of various materials such as glossiness or refraction. By simply adjusting the light with dedicated guiding vectors, we can ensure, for instance, that the intensity of highlights are maximal and minimal on both sides of the features. As with normal vectors \mathbf{n}_i for dif-

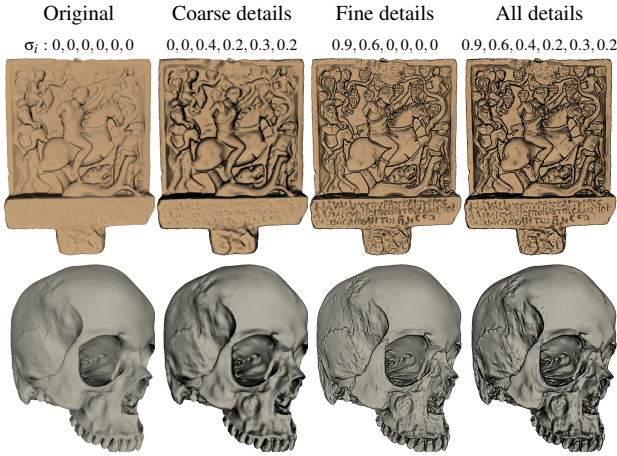


Figure 9: Diffuse enhancements of a Thracian votive relief (top) and a human skull model (bottom) using $S = 7$. Tuning properly σ_i (shown below each column title), one can easily choose which features to depict.

fuse materials, we use the reflected (respectively refracted) incident view vectors \mathbf{r}_i for specular (respectively transmissive) materials to construct the local frame $\mathcal{L}_i = (\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i)$ at scale i and compute the corresponding guiding vector \mathbf{g}_i :

$$\mathcal{L}_i : \begin{cases} \mathbf{z}_i &= \mathbf{r}_{i+1}, \\ \mathbf{x}_i &= \mathbf{r}_i - \mathbf{z}_i(\mathbf{r}_i \cdot \mathbf{z}_i), \\ \mathbf{y}_i &= \mathbf{x}_i \times \mathbf{z}_i, \end{cases}$$

$$\mathbf{g}_i = \begin{cases} \mathbf{r}_i, & \text{if } (\bar{\mathbf{l}}_{i+1} \cdot \mathbf{x}_i) \geq 0 \\ \mathcal{R}_{(\mathbf{y}_i, -\frac{\pi}{2})} \mathbf{r}_i, & \text{otherwise.} \end{cases}$$

Using these new guiding vectors we can compute a new light direction \mathbf{l}_i per shading component, allowing us to control the diffuse shading, specular highlights and refractions independently (see Figure 8). The rest of the pipeline stays unchanged.

7. Results and discussion

Results. Using our multi-scale approach, an artist can choose the proper shape depiction that suits his goal by playing with σ_i parameters. Figure 9 shows scale tuning examples for diffuse shading. Taking advantage of full shading models, any component can be used to efficiently reveal shape details. We give an example for the diffuse and specular terms in Figure 10, a): the user can choose to enhance coarse scales for the diffuse component to depict the general shape and large features of the model, while enhancing fine scales with the specular component to reveal thin details. But one could prefer to do the opposite, see Figure 10, b). In the end, the user controls exactly which scales to enhance using the appropriate shading components. This results in a wide range of possible enhancements, examples given in Figure 11. As a proof-of-concept, we also rendered global illumination results using a homemade of-fine path tracer.




Figure 10: Specular and diffuse components can be controlled independently. In a), the diffuse term is used to enhance coarse scale details with $\sigma_i = \{0., 0., 0., 0.4, 0.7, 0.6\}$ while highlights are used to depict thin features with $\sigma_i = \{0.7, 0.8, 0.3, 0., 0., 0.\}$. The opposite is done in b) by inverting σ_i .

Congruence score. To verify whether our approach improves shape depiction, we compare the shading flows obtained with and without enhancement with respect to the shape flows. As in [VBBF16], these flows are computed using structure tensors [Big87, BvdBL*06] applied per pixel in image-space. They locally specify how shape and shading patterns are stretched or compressed. More specifically, the direction of the shading flow \mathbf{v}_s is determined by extracting the maximum eigenvector of the tensor and it indicates the direction in which the compression of shading patterns is maximum. We also compute the overall magnitude of shading compression by taking the averaged eigenvalues of the structure tensor, denoted e_s . The shape flow information, similarly noted \mathbf{v}_c and e_c are obtained in the same way, but using depth gradients in the structure tensor formula instead of shading gradients.

Once the shape and shading flows are obtained, we design a *congruence* score $s \in [0, 1]$ that represents both their direction and magnitude at each pixel i of an image:

$$s = \frac{\sum_{i=1}^n e_c(i)(1 - \lfloor e_c(i) - e_s(i) \rfloor) |\mathbf{v}_c(i) \cdot \mathbf{v}_s(i)|}{\sum_{i=1}^n e_c(i)}, \quad (2)$$

where the symbol $\lfloor x \rfloor$ stands for the clamping of x in the range





Figure 11: Results gallery (see the supplementary video for more examples). The two last scenes are rendered in a global illumination context using a transmissive model [WMLT07] for that last example.

$[0, 1]$. Intuitively, this score compares the directions of flows (with the scalar product), and averages them with a weight proportional to the magnitude of the shape flow. Curved surfaces will thus have much more impact on the final score than flat ones, for which the directions become undefined. The clamping operation prevents the score from dropping in case of exaggeration (i.e. if $e_s > e_c$).

To analyze the similarities between shape and shading features on a given scene, we compute this score on a set of images rendered with 21 different lighting directions, numbered as in the inset figure. A visual comparison of the scores is given in Figure 12. Renderings obtained with our approach (third column) exhibit much more details with respect to the original renderings (first column). The corresponding scores are also higher with our approach, as shown in the false color images, where red and blue respectively correspond to low and high scores. This is especially true when the rendering is subject to masking effects, such as the one obtained with the grazing direction of light #7. In that case, the left part of the original scene appears entirely dark whereas our approach manages to reintroduce most of the shape details while still conveying a global coherent look for such a lighting direction. Comparisons with other scenes are given in the supplementary document.



Comparison with previous work. A visual comparison of our approach with the closest related works is given in Figure 14. Our method (b) reintroduces most of the details lost with simple Lambertian shading (a). The Exaggerated Shading approach [RBD06] (c) successfully depicts all shape details thanks to its modified half-Lambertian shading function that prevents any masking effects. However, such a function also tends to flatten the surface, making the main direction of light difficult to read. As a result, the perception of the shape also seems to change when the light is moving, as shown in the supplementary video.

Light Warping [VPB*09] (d) is similar to our method as the goal is to modify the directions of the lights to improve the perception of shape, by increasing the compressions of shading patterns. Yet this method may reverse dark and bright patterns on convexities or concavities and does not manage to remove masking effects caused by a lack of variations in the environment itself. Moreover, this method does not take multi-scale details into account.

Radiance Scaling [VPB*11] (e) is also limited to a single scale and directly modifies the outgoing radiance according to surface curvature, making convexities and concavities respectively darker and brighter (or inversely) using a multiplicative factor. Consequently, this method modifies the shading function itself, making the material properties look different. On regions affected by masking effects, where the resulting intensity is very low, this method also fails at reintroducing the details of the shape.

We finally compare these techniques using our congruence score, as shown in Figure 13. In this plot, the scores are computed on the renderings of the Figure 14, with the whole set of 21 different lighting directions. First, we see that all the methods produce high scores, except when the original light is positioned away from the

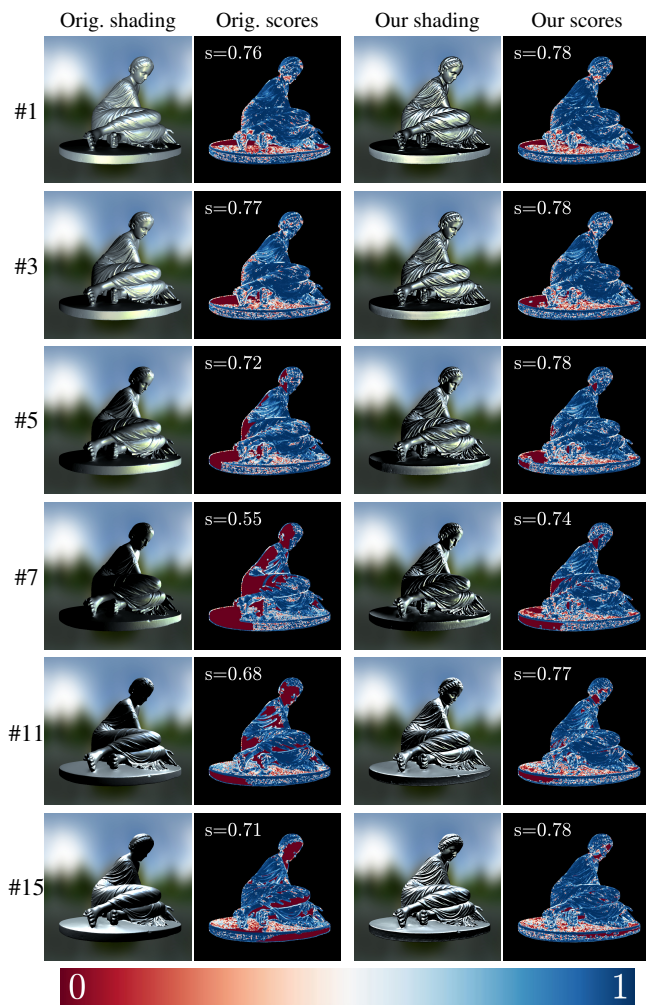


Figure 12: Renderings and congruence scores obtained without (left) and with (right) our approach. Each row corresponds to a different lighting direction, numbered as in the inset figure of Section 7. For each score image the average congruence score s is displayed in the top left corner. The colormap used to represent the per pixel score is shown at the bottom.

camera, producing a lot of masking effects. This is due to the fact that our score computation mostly compares the directions of flows rather than their magnitudes. A high score then corresponds to a good shape depiction but does not scale with the amount of exaggeration (*i.e.* when the contrast of shading patterns is increased). Overall, the score obtained with our approach is higher than all the other methods independent on the direction of light. This illustrates that, even in presence of masking effects (*i.e.* around lighting directions #7 and #21), our method succeeds at depicting shape details. Note also that Exaggerated Shading scores tend to be similar to ours. It is sometimes higher or lower depending on the input scene, as shown in the supplementary results. This is due to the fact that this method relies on a half-Lambertian shading, which ensures that the shape is visible everywhere. It is the strength of this technique,




Figure 13: Average congruence scores of previous techniques and our enhancement on the model shown in Figure 14, with 21 different lighting directions (x -axis). A Lambertian shading is used for all the techniques except Exaggerated Shading that relies on its own modified half-Lambertian shading function. The corresponding renderings can be seen in the supplementary document, along with analyses on other scenes.

but also its main limitation, as no other shading function can be used and material parameters cannot be controlled.

Limitations and future work. As our congruence score shows, our method improves the overall shape depiction in a measurable fashion with respect to previous techniques. Nevertheless, the score formula does not represent all phenomena happening in the enhanced images: exaggeration is not properly handled and the coherency with the original scene is not taken into account. Subjective evaluations could help to narrow down the impact of these properties on shape depiction. This avenue of research proves to be interesting and a possible key to better understand the relationships between shading and shape.

Pushing σ_i to extreme values can alter the original perceived light and material properties. It may be desirable in some use-cases for which a high degree of realism is not a priority. In other cases, we would like to investigate the range of σ_i values that automatically lead to plausible enhancements, *e.g.* using metrics similar to [NWH16] to guide the enhancement. New perceptual studies evaluating the HVS sensitivity to light inconsistencies within single objects will probably be needed to understand how multi-scale enhancement affects the plausibility of the scene.

We also would like to design more intuitive controls in future work as it is a key question for non-photorealistic applications. Indeed, when working with a high number of scales, a simplistic user interface made of sliders can be cumbersome to use. One could imagine better ways to interact, *e.g.* [PDF15].




Figure 14: a) Original Lambertian renderings. b) Our approach. c) Exaggerated Shading [RBD06], using $k_i \propto \sigma_i^{-0.5}$ as the default proposed in the paper. d) Light Warping [VPB*09] using $\alpha = 16.62$ (the amount of warping). e) Radiance Scaling [VPB*11] using $\lambda = 0$ (isotropic case), $\alpha = 0.5$ (what authors call the invariant point) and $\gamma = 4$ (the strength of the scaling). Each column shows a different lighting direction.

8. Conclusion

We presented an approach whose main goal is to enhance human visual perception of surface details by ensuring congruence between shading and shape flows. We achieve multiple scales shape depiction in the spirit of Exaggerated Shading [RBD06] while retaining the material diversity enabled by Light Warping [VPB*09]. We introduced a novel congruence score to objectively measure shape depiction. Our method outperforms previous work in that regard while retaining their strengths. It is general enough to be easily integrated in various rendering pipelines, from real-time to offline.

Acknowledgments

We thank Sunrise Wang and the anonymous reviewers for their useful insights and interesting comments. We also thank Arthur Novat from [Atelier Novat](#) for his time and enlightening feedback, as well as Jonathan Granier who made a prototype for single scale enhancement during his internship at Inria. Clément Baticle kindly granted us permission to use his [Romanesco scene](#). Models come either from the aim@shape collection: [bimba](#), [fertility](#), or are under the [CC Attribution licence](#): [shield](#), [Thracian bas-relief](#), [lion](#),

[skull](#), [dragon](#), [jacks player](#), [crocodile](#), [locomotive](#), [snail](#). Environment maps come from [Texturify](#).

References

- [AKAC07] ABDULLAH-AL-WADUD M., KABIR M. H., AKBER DEWAN M. A., CHAE O.: A dynamic histogram equalization for image contrast enhancement. *IEEE Transactions on Consumer Electronics* 53, 2 (2007), 593–600. 2
- [BCCS12] BANTERLE F., CORSINI M., CIGNONI P., SCOPIGNO R.: A low-memory, straightforward and fast bilateral filter through subsampling in spatial domain. *Computer Graphics Forum* 31, 1 (February 2012), 19–32. URL: <http://vcg.isti.cnr.it/Publications/2012/BCCS12>. 5
- [Big87] BIGUN J.: Optimal orientation detection of linear symmetry. In *Proceedings of the IEEE First International Conference on Computer Vision, London, Great Britain* (1987), pp. 433–438. 6
- [BSZ01] BEN-SHAHAR O., ZUCKER S.: On the perceptual organization of texture and shading flows: from a geometrical model to coherence computation. In *CVPR 2001*. (2001), vol. 1, pp. I-1048–I-1055 vol.1. 1
- [BvdBL*06] BROX T., VAN DEN BOOMGAARD R., LAUZE F., VAN DE WEIJER J., WEICKERT J., MRÁZEK P., KORNPROBST P.: *Adaptive Structure Tensors and their Applications*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006, pp. 17–47. URL: https://doi.org/10.1007/978-3-540-32867-9_2.

- 1007/3-540-31272-2_2, doi:10.1007/3-540-31272-2_2.6
- [CF07] CANIARD F., FLEMING R. W.: Distortion in 3D shape estimation with changes in illumination. In *APGV '07: Proc. symposium on Applied perception in graphics and visualization* (2007), ACM, pp. 99–105. doi:http://doi.acm.org/10.1145/1272582.1272602. 2
- [CST05] CIGNONI P., SCOPIGNO R., TARINI M.: A simple normal enhancement technique for interactive non-photorealistic renderings. *Computer & Graphics* 29, 1 (feb 2005), 125–133. URL: <http://vcg.isti.cnr.it/Publications/2005/CST05.2>
- [CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1 (1982), 7–24. URL: <https://doi.org/10.1145/357290.357293>, doi:10.1145/357290.357293. 1
- [DFRS03] DECARLO D., FINKELSTEIN A., RUSINKIEWICZ S., SANTELLA A.: Suggestive contours for conveying shape. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 22, 3 (July 2003), 848–855. 2
- [FTA04] FLEMING R. W., TORRALBA A., ADELSON E. H.: Specular reflections and the perception of shape. *J. Vis.* 4, 9 (9 2004), 798–820. URL: <http://journalofvision.org/4/9/10/>, arXiv:<http://journalofvision.org/4/9/10/Fleming-2004-jov-4-9-10.pdf>. 1
- [HCBZ01] HUGGINS P., CHEN H., BELHUMEUR P., ZUCKER S.: Finding folds: On the appearance and identification of occlusion. vol. 2, pp. 718–725. doi:10.1109/CVPR.2001.991035. 1
- [JDA07] JUDD T., DURAND F., ADELSON E. H.: Apparent ridges for line drawing. *ACM Trans. Graph.* 26, 3 (2007), 19. 2
- [KHRA*18] KUNSBURG B., HOLTSMANN-RICE D., ALEXANDER E., CHOLEWIAK S., FLEMING R., ZUCKER S.: Colour, contours, shading and shape: Flow interactions reveal anchor neighbourhoods. *Interface Focus* 8 (08 2018), 20180019. doi:10.1098/rsfs.2018.0019. 1
- [KST08] KOLOMENKIN M., SHIMSHONI I., TAL A.: Demarcating curves for shape illustration. In *ACM SIGGRAPH Asia 2008 Papers* (New York, NY, USA, 2008), SIGGRAPH Asia '08, Association for Computing Machinery. URL: <https://doi.org/10.1145/1457515.1409110>, doi:10.1145/1457515.1409110. 2
- [KvD80] KOENDERINK J., VAN DOORN A.: Photometric invariants related to solid shape. *Optica Acta: International Journal of Optics* 27, 7 (1980), 981–996. 1
- [LMLH07] LEE Y., MARKOSIAN L., LEE S., HUGHES J.: Line drawings via abstracted shading. *ACM Trans. Graph.* 26 (07 2007), 18. doi:10.1145/1276377.1276400. 2
- [LSSG10] LOPEZ-MORENO J., SUNDSTEDT V., SANGORRIN F., GUTIERREZ D.: Measuring the perception of light inconsistencies. In *Proceedings of the 7th Symposium on Applied Perception in Graphics and Visualization, APGV 2010, Los Angeles, California, USA, July 23-24, 2010* (2010), ACM, pp. 25–32. URL: <https://doi.org/10.1145/1836248.1836252>, doi:10.1145/1836248.1836252. 2
- [MA14] MOONEY S. W., ANDERSON B. L.: Specular image structure modulates the perception of three-dimensional shape. *Current Biology* 24, 22 (2014), 2737–2742. URL: <http://www.sciencedirect.com/science/article/pii/S0960982214012706>, doi:http://dx.doi.org/10.1016/j.cub.2014.09.074. 2
- [Mil94] MILLER G.: Efficient algorithms for local and global accessibility shading. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1994), SIGGRAPH '94, ACM, pp. 319–326. URL: <http://doi.acm.org/10.1145/192161.192244>, doi:10.1145/192161.192244. 2
- [MKA11] MARLOW P., KIM J., ANDERSON B. L.: The role of brightness and orientation congruence in the perception of surface gloss. *Journal of Vision* 11, 9 (08 2011), 16–16. URL: <https://dx.doi.org/10.1167/11.9.16>, arXiv:https://jov.arvojournals.org/arvo/content_public/journal/jov/932799/jov-11-9-16.pdf, doi:10.1167/11.9.16. 2
- [NWH16] NADER G., WANG K., HÉTROUY-WHEELER F., DUPONT F.: Just noticeable distortion profile for flat-shaded 3d mesh surfaces. *IEEE Transactions on Visualization and Computer Graphics* 22, 11 (2016), 2423–2436. doi:10.1109/TVCG.2015.2507578. 8
- [OCS01] OSTROVSKY Y., CAVANAGH P., SINHA P.: Perceiving Illumination Inconsistencies in Scenes. In *MIT AIM* (2001). 2
- [PDF15] PERIN C., DRAGICEVIC P., FEKETE J.-D.: Crossets: Manipulating Multiple Sliders by Crossing. In *Proceedings of Graphics Interface Conference (GI)* (Toronto, Canada, June 2015), Canadian Information Processing Society, pp. 233–240. URL: <https://hal.inria.fr/hal-01152074>. 8
- [PG04] PHARR M., GREEN S.: *GPU Gems*. Addison-Wesley, 2004, ch. Ambient Occlusion. URL: http://http.developer.nvidia.com/GPUGems/gpugems_ch17.html. 2
- [RBD06] RUSINKIEWICZ S., BURNS M., DECARLO D.: Exaggerated shading for depicting shape and detail. *ACM Trans. Graph.* 25, 3 (July 2006), 1199–1205. URL: <http://doi.acm.org/10.1145/1141911.1142015>, doi:10.1145/1141911.1142015. 2, 3, 4, 7, 9
- [RRAAW*16] RAHMAN S., RAHMAN M. M., ABDULLAH-AL-WADUD M., AL-QUADERI G. D., SHOYAIB M.: An adaptive gamma correction for image enhancement. *EURASIP Journal on Image and Video Processing* 35 (10 2016). doi:10.1186/s13640-016-0138-1. 2
- [RSI*08] RITSCHER T., SMITH K., IHRKE M., GROSCH T., MYSZKOWSKI K., SEIDEL H.-P.: 3d unsharp masking for scene coherent enhancement. *ACM Trans. Graph.* 27, 3 (Aug. 2008), 90:1–90:8. URL: <http://doi.acm.org/10.1145/1360612.1360689>, doi:10.1145/1360612.1360689. 2
- [ST90] SAITO T., TAKAHASHI T.: Comprehensible Rendering of 3-D Shapes. In *Proc. ACM SIGGRAPH '90* (1990), ACM, pp. 197–206. doi:http://doi.acm.org/10.1145/97879.97901. 2
- [Sta00] STARK J. A.: Adaptive image contrast enhancement using generalizations of histogram equalization. *IEEE Transactions on Image Processing* 9, 5 (2000), 889–896. 2
- [VBBF16] VERGNE R., BARLA P., BONNEAU G.-P., FLEMING R. W.: Flow-guided warping for image-based shape manipulation. *ACM Trans. Graph.* 35, 4 (July 2016), 93:1–93:12. URL: <http://doi.acm.org/10.1145/2897824.2925937>, doi:10.1145/2897824.2925937. 1, 6
- [VBF12] VERGNE R., BARLA P., FLEMING R., GRANIER X.: Surface Flows for Image-based Shading Design. *ACM Transactions on Graphics* 31, 4 (Aug. 2012), 94:1–94:9. URL: <https://hal.inria.fr/hal-00702280>, doi:10.1145/2185520.2185590. 1
- [VPB*09] VERGNE R., PACANOWSKI R., BARLA P., GRANIER X., SCHLICK C.: Light Warping for Enhanced Surface Depiction. *ACM Transactions on Graphics* 28, 3 (July 2009), 25:1–25:8. URL: <https://hal.inria.fr/inria-00400829>, doi:10.1145/1531326.1531331. 2, 7, 9
- [VPB*11] VERGNE R., PACANOWSKI R., BARLA P., GRANIER X., SCHLICK C.: Improving Shape Depiction under Arbitrary Rendering. *IEEE Transactions on Visualization and Computer Graphics* 17, 8 (June 2011), 1071–1081. URL: <https://hal.inria.fr/inria-00585144>, doi:10.1109/TVCG.2010.252. 2, 7, 9
- [WDK12] WIJNTJES M. W. A., DOERSCHNER K., KUCUKOGLU G., PONT S. C.: Relative flattening between velvet and matte 3d shapes: Evidence for similar shape-from-shading computations. *Journal of Vision* 12, 1 (2012), 2. doi:10.1167/12.1.2. 2
- [WMLT07] WALTER B., MARSCHNER S., LI H., TORRANCE K.: Microfacet models for refraction through rough surfaces. pp. 195–206. doi:10.2312/EGWR/EGSR07/195-206. 7