Detail Highlighting using a Shadow Edge Histogram

A. Gkaravelis[1] and G. Papaioannou[1]

[1] Athens University of Economics and Business, Greece

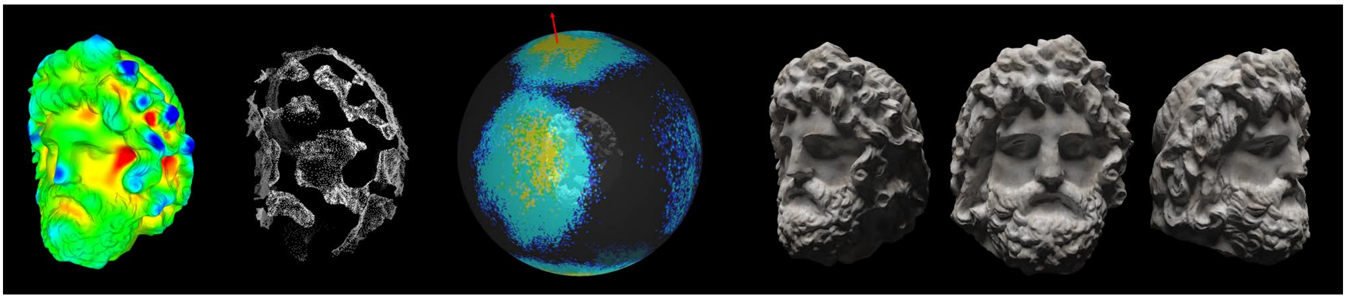

Figure 1: *Determination of a single optimal highlighting direction for the Sculpture model. From left to right: Mean curvature estimation, extraction of (recessed) interest points, histogram of shadow edge directions on interest points and resulting illumination.*

Abstract

In this paper we propose a simple and effective technique for setting up a configuration of directional light sources to accentuate the prominent geometric features of complex objects by increasing the local shadow contrast near them. Practical applications of such a task are encountered among others in professional photography, and cinematography. The method itself, which is based on a voting mechanism, quickly produces consistent and view-independent results, with minimal user intervention.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—shading, shadowing I.3.4 [Computer Graphics]: Graphics Utilities—Graphics editors

1. Introduction and Related Work

The illumination of an environment plays a very important role in establishing a specific mood, enhancing storytelling or capturing the attention of the intended audience. Choosing the appropriate lighting conditions in order to make the features of an object more pronounced and accentuate its details is not a trivial task. This is especially true for objects with many, seemingly equally important geometric features and no apparent predominate viewing angles. Although the particular process is only a narrow aspect of the general *lighting design* field, it addresses a practical problem often encountered in professional photography, exhibition setups and cinematography, where the lighting setup must be used on real objects and scenery.

In this paper we propose a novel technique for enhancing the appearance of objects by automatically setting up light sources that increase the local contrast at prominent geometric features. The method generates a histogram of potential lighting directions which cause a shadow transition between depressions and protrusions of the geometry and thus accentuate the relief of the surfaces (see Figure 1). The k prevailing directions are singled out after a clustering process and used to set up directional or physically-based luminaries at a distance from the highlighted object to emphasize its structural details.

The subject of goal-driven estimation of lighting parameters or the *inverse lighting problem* has been studied by researchers in the past, but with an emphasis on direct control over the intended effect of the light source parameterization, typically using simple methods such as painting the desired illumination level on the geometry or sketching the required intensity contrast. Most often, an optimization process attempts to minimize the difference between the *light intentions* (i.e. the goal illumination at the given points) and the light produced by a light parameterization state, including the number of sources in some cases.

Poulin and Fournier [PF92] proposed an approach to define light positions in a scene by directly manipulating their shadows or resulting highlights on both diffuse and specular surfaces. Later, they

DOI: 10.2312/egsh.20171017

extended their work and presented a system [PRJ97] that controls the position of a number of light sources by sketching strokes where the lighting effect, either shadow or highlight, should be placed. Pellacini et al. [PBMF07] proposed a similar workflow, where artists could instead directly paint the desired illumination onto the scene. A goal image was created based on the user's input and the various parameters of the light sources were optimized, in order to minimize the difference between the produced and desired images. Shacked and Lischinski [SL01] relied on the weighted sum of a number of image-based perceptual metrics as an objective function to optimize the lighting parameters that resulted in the image, using the gradient descent local optimization method.

In the context of object-centric inverse lighting problem, Pellacini [Pel10] presented an interactive interface, named *EnvyLight*, for editing natural illumination using sketching strokes on lighting features that alters regions of an environment map. In contrast, Okabe et al. [OMSI07] proposed a method that allowed users to specify the appearance of the final image by optimizing the obtained illumination from the environment map, after directly painting the incident radiance on the target model.

Lee et al. [LHV04] proposed an effective approach for the visualization of scientific datasets using non physically plausible lights; they use a surface segmentation algorithm based on curvature to split the surface into patches and optimize a light source for each one. For each segment, they assign a light direction, whose contribution to the patch minimizes a custom curvature-intensity metric, in order to emphasize the local structure of the object. Additionally, for each depth discontinuity, detected by sampling the depth buffer, they add a shadow casting direction placed at an angle between the view direction and the depth gradient. Finally, Vergne et al [VPB*10] proposed a modification of the shading of objects by scaling their incoming radiance to enhance feature depiction.

In contrast to the state of the art, we propose here an automatic method with minimal interaction from the user. Since it does not rely on an optimization process, the method is fast and simple, yet generic and practical enough to be used in a wide range of application scenarios, including lighting design for real objects. Furthermore, the illumination goal, being geometry-driven, produces view-independent results.

2. Method Overview

Our method computes the directions from which the model should be illuminated, in order to highlight its geometric details. A key observation is that to accentuate the relief patterns of a surface, local contrast at elevation transitions must be enhanced. To do so in a non-artificial manner like using ambient obscurance, a proper selection of lighting directions must be chosen so that the majority of relief transitions also coincide with accented lighting transitions. This way, the resulting illumination can be plausibly recreated in lighting setups for enhancing the appearance of real objects.

The method first isolates a set of *interest points* that reside on depressions of high curvature and therefore, with a high potential to be adjacent to protruding elements. For this task, we rely on the mean curvature estimated at sample locations on the target model.

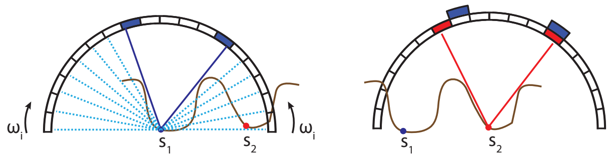

Figure 2: *Voting Mechanism. For each evaluation point ($\mathbf{s}_i$), elevations are sampled iteratively (ω_i), from the base of the hemisphere up to the apex, until we discover an open direction, which is in turn added to the shadow edge histogram (blue and red bin votes).*

In our implementation, since we use relatively densely and uniformly digitized objects, we directly compute the mean curvature on the meshes using MeshLab and export the curvature as a vertex attribute. Of course, any other computation approach or even a different metric can be employed for this task. The user decides which points will be included in the interest point set, by simply specifying a curvature threshold.

Next, for all interest points, we detect directions that form a horizon of shadow transitions or *shadow edges*. A shadow edge is formed between blocked and non-blocked parts of the hemisphere by other geometry. By requiring that potential lighting directions coincide with shadow edges, the interest points are partially obscured by the high elevation geometry, while in contrast the latter are (potentially) lit by the particular light source. The shadow edge directions add a vote to a directional histogram. The process is described in detail next, in Section 2.1.

Subsequently, the directions that most interest points voted as shadow edges are considered as candidate lighting directions. Due to the fact that similar, highly preferred directions may not correspond to exactly the same histogram cell, a clustering and cleanup follows, which results in shadow edge groups whose representative directions signify a potential lighting direction. This step is explained in Section 2.2.

2.1. Shadow Edge Histogram

For each interest point **p**, we scan the normal-oriented hemisphere from base to apex, to identify shadow edges. Rays are generated and cast from **p** using a stratified azimuthal angle and an increasing fixed altitude angle step. For each direction we cast a ray and trace it in the scene to check if the ray intersects the geometry (occluded direction). A shadow edge occurs at a transition between an occluded ray and an unoccluded one. On such event, we mark the *unoccluded* direction by incrementing its vote count in the corresponding shadow edge histogram bin. Figure 2 demonstrates this voting principle.

2.2. Selection of Lighting Directions

Depending on the particular object, the voting stage can spread the shadow edges in many similar directions, which form rough clusters in the histogram, as shown in Figure 4-middle column. We identify the dominant direction aggregations by first removing all directions which have a vote count lower than 33% of the largest

Figure 3: *Illumination of the Napoleon model using the best three clusters of directions (individually and combined).*

bin count and then applying K-means clustering, where K is the number of desired light sources. We opted for the K-means method because it provides a fast clustering and creates separable direction clusters. Separable clusters are very important because they provide better coverage of the various regions of the model. As usual in K-means clustering, each observation belongs to the cluster with the closest mean value. In our approach, we define the distance between each direction and the mean cluster direction as their dihedral angle. After the clustering process, we select as representative direction from each cluster the one with the highest voting score to be used as an illumination direction. We do not use the mean of the cluster as a representative direction because it might not correspond to a high-ranking or even valid (voted) histogram bin. In Figure 3, we can see the contribution of each one of the selected illumination directions that our method produced.

3. Results

We performed tests on various models with diverse complexity, some of which are shown in Figure 4. The first column shows the colorized curvature of the model, while the second column presents the selected interest points after thresholding the mean curvature. The resulting shadow edge histogram can be seen in the third column, where the concentration of votes in direction cells is displayed as a heat map. Finally, the last columns present the models illuminated by the computed lighting directions. In order to quantify the performance of our method, we measure for each model the percentage of shadow transitions, reported in parentheses, between the set of selected interest points and their one-ring neighboring points.

In more detail, in the *Terrain* (71%) test case we wanted to calculate a single lighting direction to achieve the same dramatic effect as that of a sunrise or sunset. The resulting lighting direction was computed at an oblique angle to the terrain, at which the features of the landscape are more visually pronounced. The *Metope* model (76%) was a similar case, where we wanted to find at least one lighting direction that could highlight the folds of the figures' garments, without washing out their details. For two lights, the method calculated directions that distributed shadows to most relief depressions and avoided interference.

Both the *Lu Yu* (61%) and the *Napoleon* (90%) models have some fine details on the body and on the creases of the clothes, which are also completely blocked by certain lighting directions. Our method properly highlights most of the protrusions and minimizes the excessive shadowing. As can be observed both here and in Figure 3, the K-means clustering ensures the proper separation of the resulting lighting directions.

Finally, the last row shows a *Fox Skull* (87%) model, which is a rather difficult case due the many cavities and respective resulting occlusion. With two lighting directions, the principal cavities, the structure of the skull as well as the teeth are well delineated.

4. Conclusion and Future Work

We presented an automatic method to generate a configuration of directional or distant light sources that enhances the appearance of a model by increasing the local contrast at prominent geometric features. Our method does not depend on an optimization process and it requires minimal user intervention. As a future research direction, we would like to explore solutions based on the shadow edge histogram also for the calculation of the position of luminaries *near* an object for the purpose of feature highlighting.

Acknowledgements

This work was partially supported by the Athens University of Economics and Business Research Centre. The Lu Yu, Fox Skull and Napoleon (©IMA Solutions and Musée de la Révolution française) models were licensed under the Creative Commons Attribution 4.0 International License by Artec Group inc. and downloaded from Sketchfab. The Scultpure model was licensed under the the Attribution-NonCommercial-ShareAlike 3.0 Unported license and downloaded from the Autodesk 123D website.

References

[LHV04] LEE C. H., HAO X., VARSHNEY A.: Light collages: Lighting design for effective visualization. In *IEEE Visualization* (2004), pp. 281 – 288. 2

[OMSI07] OKABE M., MATSUSHITA Y., SHEN L., IGARASHI T.: Illumination brush: Interactive design of all-frequency lighting. In *Computer Graphics and Applications, 2007. PG '07. 15th Pacific Conference on* (Oct 2007), pp. 171–180. 2

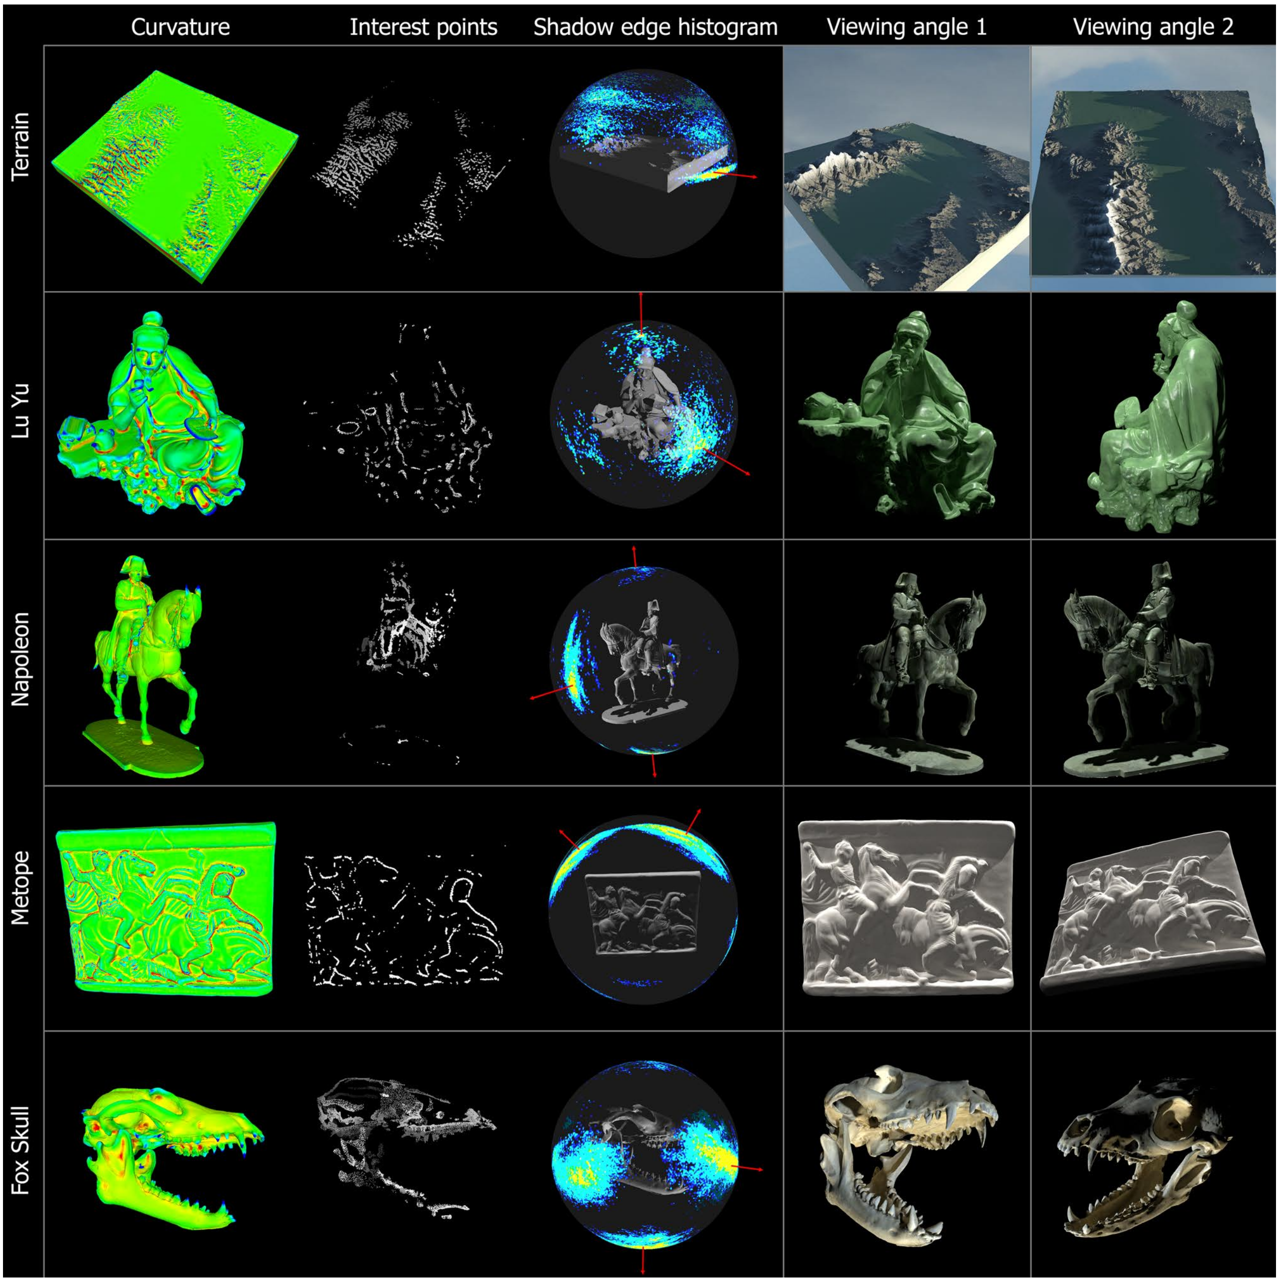

Figure 4: *Test models. Each row demonstrates from left to right, the curvature of the model, the selected interest points, the shadow edge histogram along with the selected illumination directions and the resulting illumination on each model.*

[PBMF07] PELLACINI F., BATTAGLIA F., MORLEY K., FINKELSTEIN A.: Lighting with paint. *ACM Transactions on Graphics 26*, 2 (June 2007), Article 9. 2

[Pel10] PELLACINI F.: envylight: An interface for editing natural illumination. In *ACM SIGGRAPH 2010 Papers* (New York, NY, USA, 2010), SIGGRAPH '10, ACM, pp. 34:1–34:8. 2

[PF92] POULIN P., FOURNIER A.: Lights from highlights and shadows. In *Proceedings of the 1992 Symposium on Interactive 3D Graphics* (New York, NY, USA, 1992), I3D '92, ACM, pp. 31–38. 1

[PRJ97] POULIN P., RATIB K., JACQUES M.: Sketching shadows and highlights to position lights. In *Proceedings of the 1997 Conference on Computer Graphics International* (Washington, DC, USA, 1997), CGI '97, IEEE Computer Society, pp. 56–. 2

[SL01] SHACKED R., LISCHINSKI D.: Automatic lighting design using a perceptual quality metric. *Computer Graphics Forum 20*, 3 (2001), 215–227. 2

[VPB*10] VERGNE R., PACANOWSKI R., BARLA P., GRANIER X., SCHLICK C.: Radiance scaling for versatile surface enhancement. In *Proceedings of the 2010 Symposium on Interactive 3D Graphics and Games* (New York, NY, USA, 2010), I3D '10, ACM, pp. 143–150. 2