

A Sampling Criterion for Optimizing a Surface Light Field

P. Lambert, J.-D. Deschênes, P. Hébert

Computer Vision and Systems Laboratory, Laval University, Quebec, Canada

{plambert, desche07, hebert}@gel.ulaval.ca

Abstract

This paper adopts a sampling perspective to surface light field modeling. This perspective eliminates the need of using the actual object surface in the surface light field definition. Instead, the surface ought to provide only a parameterization of the surface light field function that specifically reduces aliasing artifacts visible at rendering. To find that surface, we propose a new criterion that aims at optimizing the smoothness of the angular distribution of the light rays emanating from each point on the surface. The main advantage of this approach is to be independent of any specific reflectance model. The proposed criterion is compared to widely used criteria found in multi-view stereo and its effectiveness is validated for modeling the appearance of objects having various unknown reflectance properties using calibrated images alone.

1. Introduction

Modeling the appearance of objects is a major problem in computer vision. Here, the objective of modeling is to allow free viewpoint, photorealistic visualization of a photographed object.

To reach this objective, geometry-based approaches rely on physical principles governing the interactions between reflectance properties, light sources and geometric surfaces. By making enough assumptions concerning the reflectance properties of the object, it is possible to solve an inverse problem to obtain the shape and appearance of the object from just a few images.

The alternative image-based approaches, developed in the last decade, sample and reconstruct a light field from a dense set of viewpoints. Their major strength is the mild hypothesis under which they operate. By not assuming any specific reflectance model and being independent of the scene geometry, these approaches are well suited for visualizing objects with complex geometry and various reflectance properties. However, reconstructing a sparsely sampled light field without aliasing artifacts is still a challenge. Progress in this direction is desirable since sparse




Figure 1. From a sampling perspective, modifying the parameterization surface has geometric consequence: two rays that are indexed closely on a specific parameterization surface may not continue to be on another one, e.g. neighbouring rays 1 and 3 on the teapot surface are indexed on distant locations on the surrounding surface.

datasets reduce the complexity of both the acquisition and computation of the model.

Not surprisingly, researchers have addressed this problem by combining shape measurements with light fields [2, 3]. To benefit from the knowledge of the object shape (geometric proxy), Miller et al. [1] used a parameterization of the light field function which differs from that of the original, well-known two-plane parameterization [4, 5]. In this new parameterization, every ray of light is indexed by two indices (r, s) defining a point on the surface of the object along with two others (θ, φ) that define a direction leaving that point. Doing this leads to a so-called surface light field. Surface light fields were reported to produce sharper images than using the two-plane parameterization for comparable dataset and model sizes [2]. However, surface light field approaches forego the important advantage of light fields by relying on geometry measurements.

This paper leaves the geometric setting in favor of a sampling approach to surface light field modeling. Concretely, this implies that the (r, s) surface is no longer used to define the object surface but serves only to index the light rays, as depicted in Figure 1. The main contribution of this paper is to show how this perspective leads to a new criterion that the

surface must optimize in order to produce un-aliased rendering. Contrarily to the geometric setting, the proposed criterion is independent of any specific reflectance model yet produces models of comparable quality from images alone.

First, section 2 formalizes the problem. Then, section 3 provides an overview of the research that is related to the present work. Section 4 defines what should be the appropriate criterion to assess the optimality of a specific parameterization surface. Section 5 defines the set of potential parameterization surfaces. This requires revisiting the familiar concept of the visual hull. Section 6 finally presents an implementation and results obtained from the proposed approach.

2. Problem Statement

A surface light field provides the radiance of every ray originating from a surface. Formally, a surface light field is a 4D function:

$$I_K : K \times \mathcal{S}^2 \rightarrow \mathbb{R}^3; (r, s, \theta, \varphi) \mapsto I_K(r, s, \theta, \varphi), \quad (1)$$

where K is a mesh of the *parameterization surface*, \mathcal{S}^2 is the sphere of unit vectors in \mathbb{R}^3 and the radiance, I_K , is given by points in \mathbb{R}^3 corresponding to *RGB* triples [2]. To emphasize the implicit dependency between them, a subscript, K , of the surface light field I_K always indicates its corresponding parameterization surface.

Observe that any surface enclosing the object is a potential parameterization surface, as illustrated in Figure 1. Hence, no special importance should be given to the actual object surface; raising the following central question of this paper: *Which surface is better adapted to parameterize the surface light field and why?* An intuitive guess is that the optimal parameterization surface must be the actual surface of the object. However, justifying this not necessarily true answer is not trivial.

This problem may be formalized as a regularized optimization with constraint. The constraint arises because not every surface is an acceptable parameterization surface, e.g. a surface that encloses only a part of the object. In fact, the set κ of meshes that are potential parameterization surfaces will be given in Section 5. The optimization to perform is thus: given an initial parameterization surface $K_0 \in \kappa$ and a finite set of calibrated images of an object, here considered as a set of samples of the surface light field function, $\{I_K(r_i, s_i, \theta_i, \varphi_i)\}$, solve:

$$K_{opt} = \arg \min_{K \in \kappa} (\varepsilon_{ext}(I_K(r, s, \theta, \varphi)) + \lambda \varepsilon_{int}(K)). \quad (2)$$

K_{opt} is called the optimal parameterization surface. The external energy ε_{ext} is based on a sampling criterion applied on I_K that quantifies the optimality (to be defined later) of

the considered parameterization surface K . Finally, the internal energy ε_{int} is a function that regularizes the parameterization surface and λ is a regularization constant fixed empirically. This formalization is indeed closely related to the search for the object surface found in the multi-view setting. The purpose of this paper is however only to show that a sampling perspective imposes reviewing both the external energy ε_{ext} and the set of potential parameterization surfaces κ .

3. Related Work

When sampling a surface light field, the objective is to reconstruct the sampled function without aliasing artifacts. Aperture filtering (synthetic or not) [4] can ensure this by smoothing input images until disparity is nowhere superior to one pixel [6]. However, using this technique with a sparsely sampled light field produces excessively blurred images or may remove desirable view-dependent variations in reflectance.

The dynamic reparameterization of the light field proposed by Isaksen et al. [7] helps solve this problem. This reparameterization relies on the two-plane parameterization used in [4] and [5]. A point on an entrance plane ((s, t) coordinates) and a point on an exit plane ((u, v) coordinates) provide the four parameters required to index every ray crossing the planes. By interactively moving the exit plane through space, it is possible to control which samples in each input view will contribute to the reconstructed ray. This way, any depth may exhibit zero disparity, allowing any chosen scene element to be reconstructed without aliasing artifacts. However, artifacts will be apparent for other scene elements whose disparity exceeds one pixel. In [8] this idea was extended by allowing a user to reparameterize the light field on a series of layers. This paper generalizes these ideas by automatically computing an optimal reparameterization of the light field on a curved surface.

Our approach nonetheless requires searching for a surface, and is therefore related to the recovery of the shape found in various multi-view stereo algorithms [9]. However, the objects that are modeled here are not necessarily Lambertian and may exhibit significant view-dependent variations in reflectance. Moreover, no specific reflectance model is assumed, which makes it impossible to rely on matching across images or on a criterion based on a specific reflectance model, like that of [10, 11]. This paper is however further related to that of [10] since it also searches to optimize a surface to improve light field efficiency.

We advocate that leaving aside the necessity of using the actual object surface eases the modeling process. In this way, we are akin to the space carving approaches [12]. However, searching for the optimal parameterization surface does not assume any specific reflectance model and is




Figure 2. From 3 viewpoints, radiance along at most 3 directions is sampled at any point on the parameterization surface. On the other hand, each pixel of every image samples a point along the spatial dimensions. Typically, from every viewpoint, a camera captures millions of pixels.

therefore not directly related to the photohull.

The visual hull [13] is also a complex shape that can be obtained given a set of images alone and without assuming any reflectance model. As will be seen, this shape is however not sufficient to enable un-aliased rendering with a sparsely sampled light field. As in [14, 15], it is used as a starting point for the optimization, but it will also play a role in defining the set of all potential parameterization surfaces.

4. The Optimal Parameterization Surface

Each pixel of every image samples the radiance of a ray of the surface light field function. To reconstruct the continuous value of this function given these samples a minimum amount of smoothness between the radiance of nearby sampled rays is required. Otherwise, the interpolation of rays becomes aliased or requires a tremendous sampling rate. An optimal parameterization surface should thus increase the smoothness along the less densely sampled dimensions. The following sections explain how to achieve this goal.

4.1. The Less Densely Sampled Dimensions

In Figure 2, one can observe that the (θ, φ) dimensions of a surface light field are typically the least densely sampled. Actually, on a given parameterization surface, each pixel of each image samples a (r, s) point on this surface. Given n images each made of p pixels, at most np points are sampled along the surface (r, s) dimensions. We refer to this as the spatial sampling of the light field function. On the other hand, one viewpoint adds only one new sampled (θ, φ) direction at each point. Therefore, if one gathers n viewpoints, there is a maximum of n sampled directions per point to cover the whole hemisphere of directions leaving this point. This is called the angular sampling of the

surface light field function. Comparing the number of view-points, n , with np in typical acquisitions, it is clear that the angular sampling rate is many orders lower than the spatial sampling rate. This simple observation reveals how angular sampling is much more critical than spatial sampling.

4.2. Smoothing the Angular Dimensions

As seen in Figure 1, varying the parameterization surface definitely affects the surface light field function. To see how this has an impact on the reconstruction of the surface light field, we follow and extend the analysis of Vaish et al. [16]. This analysis is performed in flatland for simplicity. The scene is a line, located at depth d_0 , with a Lambertian surface exhibiting a black and white checkered pattern. In Figure 3(a), the hemispheres beneath the s axis at each depth d_0, d_1 and d_2 display the continuous value of the light field function for this scene, at those depths, when varying direction. In going from d_0 to d_2 , one observes that the function has less angular smoothness.

Now consider the sampling and reconstruction of the light field function at the depths d_0, d_1 and d_2 . By sampling rays from only three directions, the three cameras in Figure 3(a) would not be sufficient to correctly reconstruct the light field if the parameterization surface was either placed at d_1 or d_2 . Their continuous function is not smooth enough for only three samples to interpolate it without aliasing. On the other hand, the situation is different at d_0 since the function at this point is very smooth. Indeed, the three samples would be black and interpolation would lead to the true continuous function.

Note that if the angular sampling rate was high enough, it would be possible to reconstruct all three functions located at d_0, d_1 and d_2 . A high sampling rate would thus only ease the reconstruction process in that obtaining the optimal depth would be less critical. This corroborates the tradeoff between obtaining the exact optimal depth and increasing the sampling density [6].

Therefore, among all possible parameterizations, using one for which the angular smoothness is high should be preferred. This will allow the low angular sampling rate to nonetheless correctly reconstruct the surface light field function along these dimensions thus producing less aliased interpolated views. This suggests that a criterion be devised which specifically searches for smoothness along the angular dimension. Since the global regularity of a function depends on the decay of the amplitude of its spectrum as the frequency increases [17], we propose using, as a measure of angular smoothness, the first zero-centered moment of the amplitude spectrum along the θ dimension:

$$c_d(s) = \int |w| \left| \hat{F}_d(s, w) \right| dw, \quad (3)$$

where $\hat{F}_d(s, w) = \mathcal{F}_\theta [I_d(s, \theta)]$ is the Fourier transform

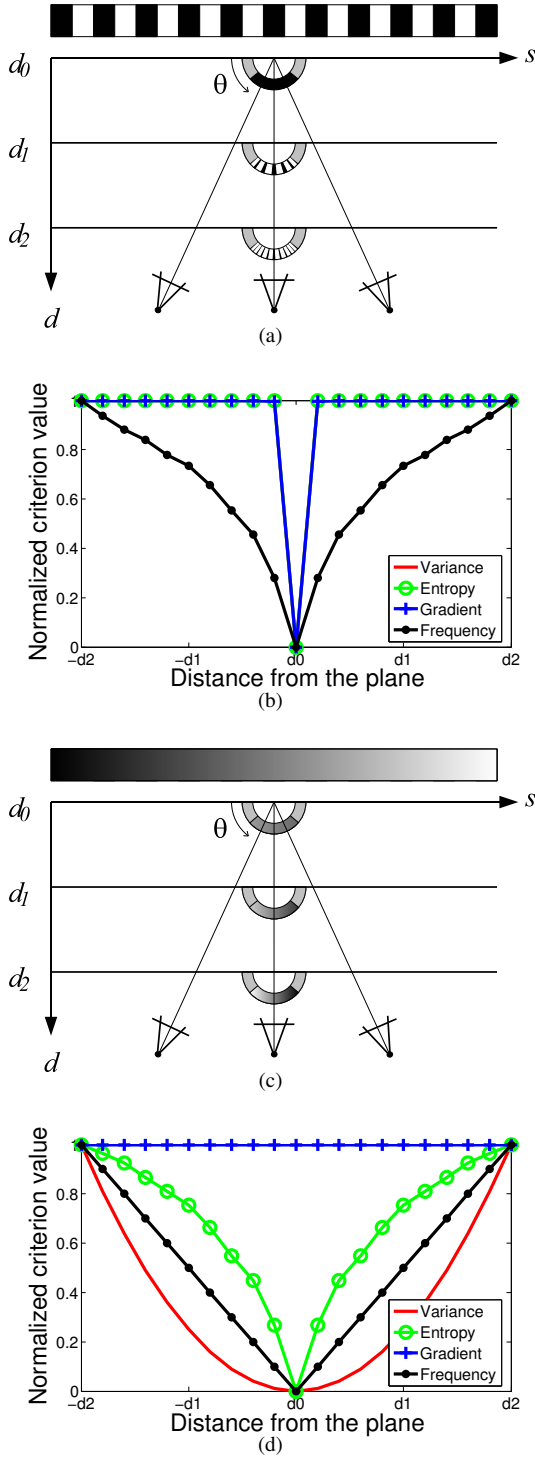


Figure 3. The consequences of changing the depth of the parameterization surface is illustrated for lines exhibiting a Lambertian checkered pattern in (a) and a constant intensity gradient in (c). The corresponding responses of various criteria are shown in (b) and (d). Note that in (b), variance, entropy and gradient criteria all have the same behaviour and are superimposed.

along the θ dimension of the surface light field function $I_d(s, \theta)$ parameterized at depth d and evaluated at point s . The equation (3) is called the *frequency criterion*. The generalization to the θ, φ domain is straightforward. Then, the external energy term ε_{ext} of equation (2) is obtained by summing the *frequency criterion* (in the θ, φ dimensions) at each of the parameterization surface point.

4.3. Comparison with previous criteria

Now let us present what makes the criterion (3) preferable over those widely used in multi-view stereo or depth from focus. Recall that the expressions of the variance v_d and gradient f_d criteria respectively used in these two fields are:

$$v_d(s) = \frac{1}{\pi} \int (I_d(s, \theta) - \bar{I}_d(s))^2 d\theta, \quad (4)$$

$$f_d(s) = - \left(\frac{\partial I_d(s)}{\partial s} \right)^2, \quad (5)$$

where $\bar{I}_d(s) = \frac{1}{\pi} \int I_d(s, \theta) d\theta$. We prefer the continuous version of these expressions to better compare them to the new proposed criterion. To further relate the *frequency criterion* to the others presented, the integrand of equation (3) may be replaced by its squared value to obtain an equivalent expression in the spatial domain as follows:

$$\hat{c}_d(s) = \frac{1}{2\pi} \int \left| \frac{\partial I_d(s, \theta)}{\partial \theta} \right|^2 d\theta, \quad (6)$$

using Parseval's identity and $iw\hat{F}_d(s, w) = \mathcal{F}_\theta \left[\frac{\partial I_d(s, \theta)}{\partial \theta} \right]$.

The evaluations of these criteria along with the entropy-based criterion of [16] are presented for the checkered pattern in Figure 3(b). When increasing depth d , we would expect the criteria to increase accordingly. However, as depicted in Figure 3(b), these criteria do not. In fact, they correctly lead to a minimum at depth d_0 , but remain constant above a certain height. In the case of the variance v_d and entropy, the explanation is that the hemispheres always display half white and half black, irrespective of their depth above d_0 . Consequently, the two criteria remain constant. The reason that the gradient criterion f_d also misbehaves is similar. The mean of every point seen from different directions soon becomes gray as the depth increases. The gradient thus also remains constant. Considering that any texture can be decomposed into a linear sum of digital sinusoid patterns of variable frequencies, each exhibiting the problem revealed, this situation may occur for a wide variety of texture patterns.

Criteria that give the same cost at d_1 and d_2 clearly miss that the function at d_1 is more accurately represented by a fixed number of angular samples than one at d_2 . This observation was the first one that led us to question the justification of variance, gradient or entropy-based criteria in

the context of surface light field rendering. These criteria are adapted to a presupposed reflectance model, when searching for the actual object surface, but they fail to put a meaningful cost to the various possible parameterizations, even for Lambertian textured objects. Stated simply, what makes the new criterion different is that it captures the angular variations of the function.

The second observation that leads us to favor this new criterion is that it applies to arbitrary reflectance models since it is based on sampling considerations instead of geometry. Contrarily to SAD, SSD or normalized cross-correlation which assume that each point has the same appearance when viewed from different directions, the criterion (3) is still suitable if significant view-dependent variations in reflectance are present. All it achieves in this case is finding a parameterization surface that displays the maximum angular smoothness. This maximum in smoothness may still exhibit significant variation along the angular dimensions.

The *frequency criterion* was also validated for a constant intensity gradient texture as in [16]. The results are shown in Figure 3(c),(d). Moreover, the results shown demonstrate that the *frequency criterion* does not have the problem of the depth from focus criterion since it correctly has a minimum at d_0 for a diffuse surface whose 2^{nd} and higher order spatial derivatives of radiance vanish.

5. The set of potential parameterization surfaces

Clearly, many surfaces are candidates for parameterizing the surface light field. Obtaining a surface that minimizes equation (2) for a given object is a complex optimization problem. It is possible to solve this problem by starting from an arbitrary surface K_0 and deforming it until the minimization is achieved, as in [18]. However, like most minimizations, starting as close as possible to the solution eases the convergence and can avoid erroneous local minima. The starting point in our case is the visual hull [13]. This starting point is coherent with our intention not to suppose any reflectance model since the visual hull is a shape that can be obtained without any such assumption [12]. Using the visual hull as a starting point is helpful but not new. Examples are the work of Hernandez and Schmitt [15] that optimizes a stereo reconstruction while enforcing a soft constraint using silhouette coherence or that of Furukawa and Ponce [14] who propose instead to rely on a hard constraint using rims.

5.1. Visual Hull as a Constraint

The true novelty in using the visual hull for surface light field modeling originates from Theorem 1 below. It provides a necessary and sufficient condition to ensure that all of the rays originating from the object and only these rays

are parameterized. Otherwise, the parameterization surface would either miss a part of the object or would incorporate a part of the background, all of which are undesirable. This condition is of great practical importance since it defines the set of potential parameterization surfaces κ and thus constrains the optimization. From a theoretical standpoint, it is interesting since it revisits the familiar concept of the visual hull and provides a new justification for it.

Theorem 1. Visual hull constraint theorem: A closed surface parameterizes all rays emerging from an object and does not parameterize any ray emerging from the background iff it has the same visual hull as that of the object.

Proof: The visual hull is defined by Laurentini [13]. His definition 2 may be restated as follows. Consider any viewpoint V outside the convex hull of the object and any point P . The external visual hull (or simply the visual hull) of an object is the set of points P for which the half-lines originating from any V and passing through P always intersect the object. A given visual hull can be identified to the set of all half-lines that arises from its definition.

A surface that parameterizes all half-lines (rays) emerging from the object and none from the background shares the same set of half-lines as that of the visual hull of the object. Thus this parameterization surface has the same visual hull as that of the object.

Let us prove the converse statement. If the visual hull of the parameterization surface includes a point that does not belong to the visual hull of the object, then it is possible to find a viewpoint for which this point projects to a background pixel. Therefore, this surface parameterizes a background ray. On the other hand, if the visual hull of the parameterization surface does not include a point of the visual hull of the object, then a viewpoint exists for which the projection of this point is an object pixel and is not included in the parameterization surface. Therefore, this parameterization surface misses an object ray. The converse statement is proved: a closed surface that has the same visual hull as that of the object parameterizes all rays of it and none of the background. \square

Since the optimization starts from the visual hull of the object itself, there is a simple way to enforce the *visual hull constraint*. Consider a point of the visual hull having at least one convex principal curvature [19] and the object ray tangent at this point and parallel to its convex principal direction. Deforming the visual hull surface by moving such a point inward automatically leads to a shape that does not have the same visual hull as that of the object since the considered ray would be missing. Therefore, the optimization of the equation (2) under the *visual hull constraint* could never move such a point. On the other hand, any other point is free to move inward. Note that moving any point of the visual hull outward never respects the visual hull constraint and is thus never allowed. Following this guideline ensures

that only a set of parameterization surfaces κ that satisfies the *visual hull constraint* is visited during the optimization and thus the constraint is enforced.

6. Implementation and results

Many multi-view stereo algorithms could be used to perform the optimization of equation (2) [9]. Since the purpose of this paper is not to delve into the intricacies of these methods, we have implemented a greedy line-search optimization, see Algorithm 1, that is simple to implement and sufficient to illustrate the validity of the proposed approach. As in the multi-view setting, the minimization to be achieved is complicated by the visibility problem [20]. This problem itself is very complex and will not be addressed here. Visibility is taken into account simply by conceiving the parameterization surface as opaque.

Algorithm 1 Greedy line-search optimization

```

 $\alpha_{max} \leftarrow$  maximum displacement distance
 $k \leftarrow$  number of steps in line-search
 $T_{curvature} \leftarrow$  maximal convex principal curvature
 $T_{Gain} \leftarrow$  minimum worth moving gain
 $K_0 \leftarrow$  visual hull mesh
 $\gamma_1, \gamma_2 \leftarrow$  principal curvatures of  $K_0$ 
 $S \leftarrow$  vertices of  $K_0$  with  $\min(\gamma_1, \gamma_2) \geq T_{curvature}$ 
for every vertex  $\in S$  located at  $\vec{q}_0$  with normal  $\vec{N}$  do
  for  $j = 0$  to  $k$  do
     $M[j] \leftarrow \hat{c}(\vec{q}_0 + j \cdot (\alpha_{max}/k) \cdot \vec{N})$ 
  end for
   $j_{min} \leftarrow \operatorname{argmin}(M); Gain \leftarrow M[0] - M[j_{min}]$ 
  if  $Gain > T_{Gain}$  then
     $\vec{q} \leftarrow \vec{q}_0 + j_{min} \cdot (\alpha_{max}/k) \cdot \vec{N}$ 
  end if
end for
Apply a smoothing operator

```

Note that the expression (6) is preferred for its computational efficiency and is evaluated using interpolation on the acquired samples $\{I_K(r_i, s_i, \theta_i, \varphi_i)\}$. Regarding this expression, only a limited cone around the normal direction is used as the domain of integration since aliasing becomes predominant at grazing angles. A cone of 60 degrees is typically used.

Both synthetic and real image sets were processed in the experiments. The acquisition consisted in capturing or simulating several calibrated images of an object against a known and fixed colored background. In real image sets, positioning was made possible by using targets laid around the objects. The images were acquired using a 6.3 megapixel Canon EOS Rebel Digital and are roughly uniformly distributed on a hemisphere around the object.




Figure 4. Results for the Buddha data set [3]. (a) One of the input pictures. (b) High resolution close up rendering of the model for visual comparison. Light Field Maps [3] using (c) the explicit geometry and (d) the optimal parameterization surface.

After positioning and extracting the silhouettes, the visual hull was reconstructed using a volumetric grid of 150^3 grid points. The proposed optimization does not depend on a particular surface light field representation. We chose the Light Field Mapping algorithm [3] to represent our surface light fields.

The first experiment was conducted on a synthetic object exhibiting complex, non-Lambertian reflectance properties. It illustrates a case for which matching across images would not be legitimate. The Buddha model from [3] was used and results obtained from its 281 pictures data set are presented in Figure 4. Note that the high number of images is necessary to capture the complex reflectance of the model. The proposed optimization leads to a parameterization that produces results that are almost identical to those of [3], shown in Figure 4(c), although the latter uses the explicit geometric model.

To further illustrate the difference between the frequency and the variance criteria, Figures 5(a),(b) show a plot of the value taken by these respective criteria as a point on the chest of the Buddha model is moved inwardly. The Buddha surface is known to lie close to the depth value 5 on the plot (pointed by the arrow in Figure 5 (a)) and the scale of the abscissa is such that its end value lies at the center of the model. Contrarily to what would be expected, the variance reaches its minimum value far from the actual object surface. The same criteria are also evaluated for neighbouring points and the result is shown in Figures 5(c)-(f), which shows that the minima are close to those previously

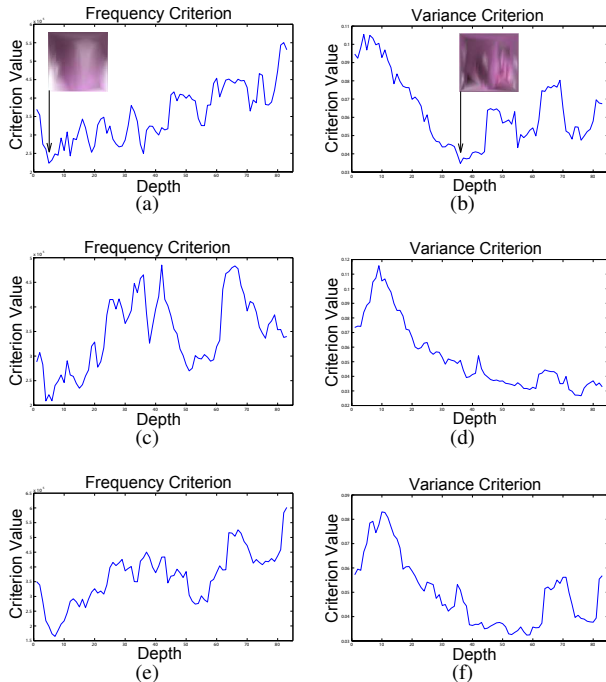


Figure 5. Frequency (a) and variance (b) criteria evaluation for a point on the Buddha according to its depth. The small images in the figures display a planar projection of surface light field values along the θ, φ dimensions at the depth corresponding to the minimum values of respective criteria. (c),(e) Frequency criterion for two other neighbouring points. (d),(f) Variance criterion for the same two points.

found. This confirms that the chosen point is not an outlier and smoothing the surface may not be appropriate to solve this problem. The reason why the variance misbehaves is the presence of a strong white specularity at this point of the surface, visible in the small image of Figure 5(a), that increases the variance of this point when compared to that visible in the small image of Figure 5(b).

Figure 6 presents results obtained from a real data set of 132 pictures. The object used is now a glossy plate laid down in a cylinder box. Parameterizing the surface light field on the visual hull (see Figures 6(a),(c)) and on the optimal parameterization surface (see Figures 6(b),(d)) is compared. Recall that comparing to the visual hull is important in this context since it is another complex shape that can be obtained without assuming any reflectance model. The concavity makes the visual hull a poor sampling surface. In spite of the glossy finish, the proposed algorithm found a parameterization surface that exhibits much less aliasing artifacts than the visual hull or any other arbitrarily chosen surface. Note that aliasing artifacts are better visualized in




Figure 6. Light field map rendering of a real data set using the visual hull (a),(c) and using the optimal parameterization surface (b),(d).

the video located on the website of the authors¹ because they have a dynamic structure.

The last series of results presented in Figure 7 is meant to display a case that still has specularities and a concavity but that now also encompasses hard edges. The object is a plastic LEGO box filled with blocks. Hard edges are a challenge for most surface reconstruction algorithms that enforce a smooth reconstruction. In spite of this, the optimal parameterization surface produces good rendering quality by adequately leveraging on the geometry/images tradeoff. This is why the model displays a realistic appearance although the optimal parameterization surface found is not the actual object surface, as seen in Figure 7(c).

7. Conclusion

Instead of relying on physical principles governing the interactions between reflectance properties, light sources and geometric surfaces, a new criterion that minimizes the first zero-centered moment of the amplitude spectrum along the angular dimensions was proposed for selecting a surface light field parameterization surface. Results showed the effectiveness of the criterion for objects with various reflectance properties. It would be interesting to further investigate the relation between the optimal parameterization surface and the actual object surface.

Acknowledgements

The authors express their gratitude to the NATEQ/FQRNT and NSERC/CRSNG for supporting this research and Annette Schwerdtfeger for proof reading.

¹<http://vision.gel.ulaval.ca/~vision3d>




Figure 7. Using the visual hull (a) as a parameterization surface produces the model shown in (b). The optimal parameterization surface (c) leads to the model shown in (d).

References

- [1] G. Miller, S. Rubin, D. Ponceleon, "Lazy Decompression of Surface Light Fields for Precomputed Global Illumination," in Proc. 9th Eurographics Workshop on Rendering, June 1998, pp. 281-292.
- [2] D. Wood, D. Azuma, K. Aldinger, B. Curless, T. Duchamp, D. Salesin and W. Stuetzle, "Surface light fields for 3D photography," in Computer Graphics (Proc. SIGGRAPH00), Aug. 2000, pp. 287-296.
- [3] W.-C. Chen, J.-Y. Bouguet, M. H. Chu and R. Grzeszczuk, "Light Field Mapping: Efficient Representation and Hardware Rendering of Surface Light Fields," ACM Transactions on Graphics, Vol. 21, No. 3, 2002, pp. 447-456.
- [4] M. Levoy and P. Hanrahan, "Light Field Rendering," in Computer Graphics (Proc. SIGGRAPH96), Aug. 1996, pp. 31-42.
- [5] S. J. Gortler, R. Grzeszczuk, R. Szeliski and M. F. Cohen, "The lumigraph," in Computer Graphics (Proc. SIGGRAPH96), Aug. 1996, pp. 43-54.
- [6] J.-X. Chai, X. Tong, S.-C. Chan, and H.-Y. Shum, "Plenoptic sampling," in Computer Graphics (Proc. SIGGRAPH00), Aug. 2000, pp. 307-318.
- [7] A. Isaksen, L. McMillan and S. J. Gortler, "Dynamically Reparameterized Light Fields," in Computer Graphics (Proc. SIGGRAPH00), Aug. 2000, pp.297-306.
- [8] H.-Y. Shum, J. Sun, S. Yamazaki, Y. Li and C.-K. Tang, "Pop-Up LightField," ACM Trans. on Graphics, Vol. 23, No. 2, April 2004.
- [9] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein and R. Szeliski, "A Comparison and Evaluation of Multi-View Stereo Reconstruction Algorithms," in Proc. IEEE Conf. Computer Vision and Pattern Recognition, Vol. 1, June 2006, pp. 519-526.
- [10] P. Ramanathan, E. Steinbach, P. Eisert and B. Girod, "Geometry Refinement for Light Field Compression," Proc. Int. Conf. Image Processing, vol. 2, Sept. 2002, pp. 225-228.
- [11] S. Soatto and A.J. Yezzi, "Tales of Shape and Radiance in Multi-view Stereo," Tech. Report CSD-02003.
- [12] K. Kutulakos and S. Seitz, "A theory of shape by space carving," in Proc. Int. Conf. Computer Vision, Vol. 1, Sept. 1999, pp. 307-314.
- [13] A. Laurentini, "The Visual Hull Concept for Silhouette-Based Image Understanding", IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. 16, No. 2, Feb. 1994.
- [14] Y. Furukawa and J. Ponce, "Carved Visual Hulls for Image-Based Modeling," in Proc. European Conference on Computer Vision, May 2006.
- [15] C. Hernández and F. Schmitt, "Silhouette and Stereo Fusion for 3D Object Modeling," in Computer Vision and Image Understanding, Vol. 96, No. 3, Dec. 2004, pp. 367-392.
- [16] V. Vaish, M. Levoy, R. Szeliski, C. L. Zitnick and S. B. Kang, "Reconstructing Occluded Surfaces using Synthetic Apertures: Stereo, Focus and Robust Measures", in Proc. IEEE Conf. Computer Vision and Pattern Recognition, Vol. 2, June 2006, pp. 2331-2338.
- [17] S. Mallat, "A Wavelet Tour of Signal Processing," Academic Press, Second Edition, 1999.
- [18] O. Faugeras and R. Keriven, "Variational principles, Surface Evolution, PDE's, Level Set Methods, and the Stereo Problem," in IEEE Trans. On Image Processing, Vol. 7, No. 3, March 1998.
- [19] M. Do Carmo, "Differential Geometry of Curves and Surfaces", Prentice Hall, 1976.
- [20] F. Durand, G. Drettakis and C. Puech, "The 3D Visibility Complex", ACM Trans. on Graphics, Vol. 21, No. 2, April 2002.