

Portal-Masked Environment Map Sampling

Benedikt Bitterli^{1,2}

Jan Novák²

Wojciech Jarosz²

¹Walt Disney Animation Studios

²Disney Research Zürich

Abstract

We present a technique to efficiently importance sample distant, all-frequency illumination in indoor scenes. Standard environment sampling is inefficient in such cases since the distant lighting is typically only visible through small openings (e.g. windows). This visibility is often addressed by manually placing a portal around each window to direct samples towards the openings; however, uniformly sampling the portal (its area or solid angle) disregards the possibly high frequency environment map. We propose a new portal importance sampling technique which takes into account both the environment map and its visibility through the portal, drawing samples proportional to the product of the two. To make this practical, we propose a novel, portal-rectified reparametrization of the environment map with the key property that the visible region induced by a rectangular portal projects to an axis-aligned rectangle. This allows us to sample according to the desired product distribution at an arbitrary shading location using a single (precomputed) summed-area table per portal. Our technique is unbiased, relevant to many renderers, and can also be applied to rectangular light sources with directional emission profiles, enabling efficient rendering of non-diffuse light sources with soft shadows.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

1. Introduction & Previous Work

Captured, image-based representations of illumination are a popular and effective technique to incorporate realistic lighting in virtual scenes. Blinn, Newell and Greene [BN76, Gre86] pioneered the use of environment maps in computer graphics, and their work was later extended to high dynamic range imagery by Debevec [Deb98]. Since then, environment lighting has remained a central tool in architectural visualization and film production for expressing rich, real-world illumination which would otherwise be difficult to represent synthetically.

In this context, we are concerned with computing the radiance L_r reflected from a shade point \mathbf{x} in direction $\vec{\omega}_o$:

$$L_r(\mathbf{x}, \vec{\omega}_o) = \int_{\Omega^+} \rho(\mathbf{x}, \vec{\omega}_o, \vec{\omega}_i) L_{\text{env}}(\vec{\omega}_i) V(\mathbf{x}, \vec{\omega}_i) d\vec{\omega}_i^{\perp}, \quad (1)$$

where ρ is the BRDF, L_{env} is the distant lighting (defined by an environment map), V is the visibility to the environment, and $d\vec{\omega}_i^{\perp}$ represents the projected solid angle measure over the hemispherical domain Ω^+ . We are interested in the efficient Monte Carlo approximation of this integral. Many techniques exist for importance sampling just the environment map [ARBJ03, KK03, ODJ04, PH10], just

the BRDF [PH10, Hd14], their product [BGH05, CJAMJ05, CETC06, SA07, RCL*08, JCI09], or some linear combination using multiple importance sampling (MIS) [Vea98].

In this paper, we are particularly interested in rendering *indoor* scenes. These are typically characterized by closed off rooms connected to the outside only through a small set of visibility openings, e.g. windows or doors. Most of the environment is therefore occluded, and the aforementioned approaches—that importance sample only the environment map and/or the BRDF while excluding visibility—become ineffective. While some of the general product sampling techniques could theoretically be extended to account for visibility, e.g. by precomputing it or “rasterizing” it to a suitable form on-the-fly, and several advanced rendering algorithms for factoring in visibility exist [Jen95, CAM08, GKPS12, VKŠ*14], they all require either significant memory usage, computation time, or complex data structures as visibility is typically unknown in general scenes.

To remedy this issue, sample guides called *portals* are frequently used in production to improve sampling efficiency. Portals are (manually) placed over windows and doors to mark connections to the outside, and samples are directed




Figure 1: An indoor scene (a) with three example shading points and the corresponding regions of the environment map visible through the window (portal) on the right wall. An environment map parametrized using global spherical coordinates (b) leads to warped visible regions, which are difficult to sample. Performing MIS between portal and environment map sampling (c) leads to a linear combination of L_{env} and V . By using our rectification (d), the visible regions become rectangular and efficient to sample.

towards the portals in hopes of hitting the environment map. Rectangular portals—considered in this paper—are most common as they balance sampling efficiency and the ability to approximate conventional architectural openings.

One would ideally like to draw light samples proportional, at each shade point, to the region of the environment map visible through the portal. Figure 1a,b shows three such *visible regions* seen from three different shade points. While the portal is rectangular, the visible regions are *warped* due to the environment map projection, and their location and shape therefore differs for each shade point. This complexity has so far made direct, and efficient, sampling of the visible environment map an elusive goal.

A practical, but imperfect, way to account for visibility is to importance sample also the area or the solid angle [Arv95, UFK13] of the portal and combine it with one of the aforementioned techniques using MIS [Vea98]. While this is lightweight and is the current best practice, the resulting probability density function is only a linear combination of the individual, importance-sampled terms (illustrated in Figure 1c), which can be quite far from the desired product (cf. Figure 1b), wasting sampling effort.

Contributions & Overview. In this paper, we propose a practical way to directly importance sample the product of the portal visibility and environment map, effectively drawing samples at each shade point proportional to the *visible environment map*. To overcome the aforementioned challenges, we observe that the shape of the visible region depends heavily on the parametrization of the environment map (Figure 1b and c use standard spherical coordinates). This observation leads us to a novel, portal-rectified, hemispherical reparametrization (Section 2) of the environment map with the critical property that the visible region from *any* shade point always maps to an axis-aligned rectangular region in the environment map. This rectified projection allows us to develop a practical algorithm to importance sample the visible environment map (Section 3) while requiring just a single, precomputed, summed-area table per portal. The resulting sampling algorithm is unbiased and inexpensive.

Furthermore, the proposed parametrization enables evaluating the total energy of the environment map visible through the portal. In scenes with several portals, we can thus choose the portal to be sampled proportionally to its expected visible contribution. We incorporate our sampling technique into a path tracer and demonstrate its performance over existing approaches in Section 4. Finally, we show that our approach can also be applied to the dual problem of sampling a rectangular area light source with a directional emission profile.

2. Rectifying the Environment Map

We propose a new parametrization of the environment map based on *rectified coordinates* (α, β) to ensure that the portal always projects onto the environment map as an axis-aligned rectangle. After resampling the environment map using rectified coordinates, we show how to efficiently importance sample the rectangular visible region in the next section.

To simplify the derivation, we express all points and directions in a canonical coordinate frame illustrated in Figure 2. The frame is centered at the shading point \mathbf{x} , rotated to align the x and y axes with the horizontal edge \vec{e}_x and vertical edge \vec{e}_y of the portal, and uniformly scaled such that the portal is contained within the canonical portal plane at $z = 1$. Note that these transformations are merely to simplify notation and do not change the visible region.

Since the edges of the portal are aligned with the coordinate axes, we can trivially enforce axis-aligned rectangular visible regions by parametrizing the environment map using the x and y axes. This is however impractical for resampling the environment map as the axes extend to infinity. We overcome this problem by expressing the horizontal and vertical positions using the following angles, illustrated in Figure 2:

$$(\alpha, \beta) = \phi(x, y) = (\arctan(x), \arctan(y)). \quad (2)$$

The inverse mapping:

$$(x, y) = \phi^{-1}(\alpha, \beta) = (\tan(\alpha), \tan(\beta)), \quad (3)$$

yields back the spatial coordinates, which can be transformed




Figure 2: A portal (blue) in the canonical reference frame centered on the shade point \mathbf{x} . In rectified coordinates, points (x, y) on the portal are parametrized using the angles (α, β) subtended with the coordinate axes.

into a direction $\vec{\omega}$:

$$\vec{\omega} = \frac{(x, y, 1)}{d}, \quad d = \sqrt{x^2 + y^2 + 1}. \quad (4)$$

This direction can then be used for evaluating the environment map for a given coordinate frame and (α, β) . We note that, although the canonical reference frame is defined by a shade point as well as the portal position and orientation, the direction $\vec{\omega}$ only depends on the portal axes \vec{e}_x and \vec{e}_y and the rectified coordinates (α, β) . This means that the environment map has to be resampled only once for each unique portal orientation, and not for each possible combination of portal and shade point, making our approach practical. An example of an environment map reparametrized in rectified coordinates is shown in Figure 3.

Jacobian Determinant. Generating samples in rectified coordinates is convenient, but integration is typically performed with respect to the solid angle measure. In order to correctly transform probability densities with respect to one measure to the other, we need to compute the Jacobian determinant of the mapping between the two spaces. We derive this determinant in two steps: first, we compute the determinant with respect to the area measure on the canonical portal plane, after which we derive the determinant with respect to solid angle.

The Cartesian x and y coordinates serve as a surface area parametrization of the canonical portal plane. The relation to rectified coordinates can be expressed as:

$$\frac{dx dy}{d\alpha d\beta} = \begin{vmatrix} \frac{\partial \phi_x^{-1}}{\partial \alpha} & \frac{\partial \phi_y^{-1}}{\partial \alpha} \\ \frac{\partial \phi_x^{-1}}{\partial \beta} & \frac{\partial \phi_y^{-1}}{\partial \beta} \end{vmatrix} = \begin{vmatrix} \sec^2(\alpha) & 0 \\ 0 & \sec^2(\beta) \end{vmatrix} = \sec^2(\alpha) \sec^2(\beta). \quad (5)$$

The relation with respect to solid angle involves multiplying by the standard geometry term

$$\frac{d\vec{\omega}}{d\alpha d\beta} = \frac{d\vec{\omega}}{dx dy} \frac{dx dy}{d\alpha d\beta} = \frac{\vec{\omega}_z}{d^2} \sec^2(\alpha) \sec^2(\beta), \quad (6)$$




Figure 3: An illustration of rectified coordinates (α, β) over the hemisphere (left), an environment map resampled using rectified coordinates (center), and a contour plot of the resulting Jacobian in Equation (7) (right).

which can be further simplified with the trigonometric identity $\sec^2(\alpha) = 1 + \tan^2(\alpha) = 1 + x^2$ and the definition of $\vec{\omega}$ to obtain:

$$\begin{aligned} \frac{d\vec{\omega}}{d\alpha d\beta} &= \frac{1}{d^3} (1 + x^2)(1 + y^2) \\ &= \frac{d^4}{d^3} \left(\frac{1}{d^2} + \frac{x^2}{d^2} \right) \left(\frac{1}{d^2} + \frac{y^2}{d^2} \right) \\ &= d (\vec{\omega}_z^2 + \vec{\omega}_x^2) (\vec{\omega}_z^2 + \vec{\omega}_y^2) \\ &= \frac{(1 - \vec{\omega}_y^2)(1 - \vec{\omega}_x^2)}{\vec{\omega}_z}. \end{aligned} \quad (7)$$

This relation only depends on the sampled direction $\vec{\omega}$ and can be readily computed during sampling. A contour plot of the Jacobian is shown in Figure 3.

Importance Map. Rectified coordinates will serve as a convenient sampling domain because they preserve axis-aligned straight lines (in contrast to spherical coordinates) and the joint distribution $p(\alpha, \beta) \propto L_{\text{env}}(\vec{\omega})$ can be decomposed into a product of marginal and conditional distributions, precomputed once for all shading points (as shown in the next section). To construct these distributions, we resample the luminance of L_{env} using the rectified coordinates into a scalar importance table I . We additionally scale the luminance by the determinant in Equation (7) to account for non-uniform stretching of the importance map due to the rectification.

Since the canonical frame (and thus ϕ) is different for each portal orientation, a separate importance map is associated with each portal. It is also worth noting that as long as we look at the portal from one side only, it is sufficient to store only the upper hemisphere of I .

3. A Practical Sampling Algorithm

In this section, we describe a practical approach to importance sample the visible environment map in rectified coordinates. Given a shading point and a portal, we can compute the visible region of the importance map $[\alpha_0, \alpha_1] \times [\beta_0, \beta_1]$ as:

$$(\alpha_0, \beta_0) = \phi(x_0, y_0), \quad \text{and} \quad (\alpha_1, \beta_1) = \phi(x_1, y_1), \quad (8)$$

where (x_0, y_0) and (x_1, y_1) are the canonical coordinates of the bottom left and top right corner of the portal (see Figure 2).

Our goal is to choose a sample point (α, β) inside the visible region with probability $p(\alpha, \beta) \propto I(\alpha, \beta)$.

We draw samples following the standard importance sampling of discrete 2D distributions [PH10]: we view the values in the visible region of I as an unnormalized, piecewise constant PDF $p(\alpha, \beta)$, from which we derive a marginalized density $p_\alpha(\alpha)$ and a conditional density $p_\beta(\beta|\alpha)$.

The main difference between our approach and traditional sampling of discrete distributions is that in our case, only a rectangular subset of I is used. This subset varies depending on the location of the shading point, meaning that the marginalized PDF is different for each \mathbf{x} . In addition, since the extents of the visible region do not generally fall on integer coordinates in the discrete table, we need to properly consider fractional table cells.

To handle these problems, we transform the discrete table to a summed-area table [Cro84], which allows evaluating integrals over rectangular windows in constant time. Linearly interpolating the summed-area table also allows us to compute integrals of fractional rectangles. We refer to the summed-area table version of I as $S(\alpha, \beta)$:

$$S(\alpha, \beta) = \int_0^\alpha \int_0^\beta I(\tilde{\alpha}, \tilde{\beta}) d\tilde{\beta} d\tilde{\alpha}. \quad (9)$$

We also define a shorthand form $R(a_0, b_0, a_1, b_1) = S(a_0, b_0) + S(a_1, b_1) - S(a_1, b_0) - S(a_0, b_1)$, which evaluates the integral of values inside a rectangle $[a_0, a_1] \times [b_0, b_1]$.

Given the visible region $[\alpha_0, \alpha_1] \times [\beta_0, \beta_1]$ induced by the shading point \mathbf{x} , we can now define the marginal cumulative distribution function (CDF):

$$P_\alpha(\alpha) = \frac{R(\alpha_0, \beta_0, \alpha, \beta_1)}{R(\alpha_0, \beta_0, \alpha_1, \beta_1)}, \quad (10)$$

and the conditional CDF:

$$P_\beta(\beta|\alpha) = \frac{R(\lfloor \alpha \rfloor_h, \beta_0, \lceil \alpha \rceil_h, \beta)}{R(\lfloor \alpha \rfloor_h, \beta_0, \lceil \alpha \rceil_h, \beta_1)}, \quad (11)$$

where h is the side length of a table cell and $\lfloor \cdot \rfloor_h$ and $\lceil \cdot \rceil_h$ round to the closest smaller and larger multiple of h , respectively. For a table with $N \times N$ grid cells, the cell side length is π/N .

Given two random variables $\xi_1, \xi_2 \in [0, 1]$, we can compute the corresponding α using $P_\alpha^{-1}(\xi_1)$ and β using $P_\beta^{-1}(\xi_2|\alpha)$. Both the marginal and conditional CDFs are piecewise linear and can be inverted exactly. We provide pseudo-code for inverting P_α in Algorithm 1. The inversion procedure follows standard bisection: first, a conservative interval of grid cells is computed that is guaranteed to contain $P_\alpha^{-1}(\xi)$. This interval is then repeatedly shrunk until it is precisely one grid cell wide. The per-grid-cell CDF is linear and can be inverted directly. The inversion method for P_β follows analogously, only using β_0, β_1 and $P_\beta(\cdot|\alpha)$ instead of the α counterparts.

For purposes of integrating over the hemisphere, the sam-

Algorithm 1: $P_\alpha^{-1}(\xi)$

```

1 lowerBound  $\leftarrow \lfloor \alpha_0 \rfloor_h$ ;
2 upperBound  $\leftarrow \lceil \alpha_1 \rceil_h$ ;
3 while upperBound - lowerBound  $\neq h$  do
4   midPoint  $\leftarrow \lfloor \frac{1}{2}(\text{lowerBound} + \text{upperBound}) \rfloor_h$ ;
5   if  $P_\alpha(\text{midPoint}) > \xi$  then
6     upperBound  $\leftarrow$  midPoint;
7   else
8     lowerBound  $\leftarrow$  midPoint;
9 return lowerBound +  $h \frac{\xi - P_\alpha(\text{lowerBound})}{P_\alpha(\text{upperBound}) - P_\alpha(\text{lowerBound})}$ 

```

pling PDF can be expressed in solid angle measure as:

$$p_{\vec{\omega}}(\vec{\omega}) = \left| \frac{d\vec{\omega}}{d\alpha d\beta} \right|^{-1} p_\alpha(\alpha) p_\beta(\beta|\alpha) = \frac{\vec{\omega}_z p_\alpha(\alpha) p_\beta(\beta|\alpha)}{(1 - \vec{\omega}_y^2)(1 - \vec{\omega}_x^2)}. \quad (12)$$

Note that since we pre-multiplied I by the Jacobian in Equation (7), its inverse in Equation (12) will approximately cancel out (up to discretization error), leading to nearly perfect importance sampling of the visible environment map.

4. Implementation and Results

We integrated our importance sampling into a custom path tracer as an alternative light sampling strategy for portals and rectangular lights with directional emission profiles (specified in IES format [III91]). Upon loading a scene, we precompute and store the summed-area table S for each portal/rect light. Note that portals with parallel axes, e.g. pairs of windows on the same wall in the HORSE ROOM scene, can share the same S . During rendering, we estimate direct illumination using one BRDF sample and one light sample, which we combine using the power heuristic [Vea98] with exponent two.

Probabilistic Portal Selection. When multiple portals/lights are present, we would like to distribute the samples among the portals based on each portal's expected contribution. A straightforward approach is to approximate the contribution by the subtended solid angle of the portal. An added benefit of our rectification, however, is that the summed-area table enables integrating (in constant time) the total energy of the environment map visible through the portal from any shade point. We can therefore probabilistically select the portals for sampling proportionally to their total expected energy, thereby further reducing the variance. We refer to this as *portal-visible energy* selection.

4.1. Evaluation

We evaluate the benefits of our approach on two indoor scenes with distant environment lighting, and an indoor scene with rectangular lights with directional emission profiles.

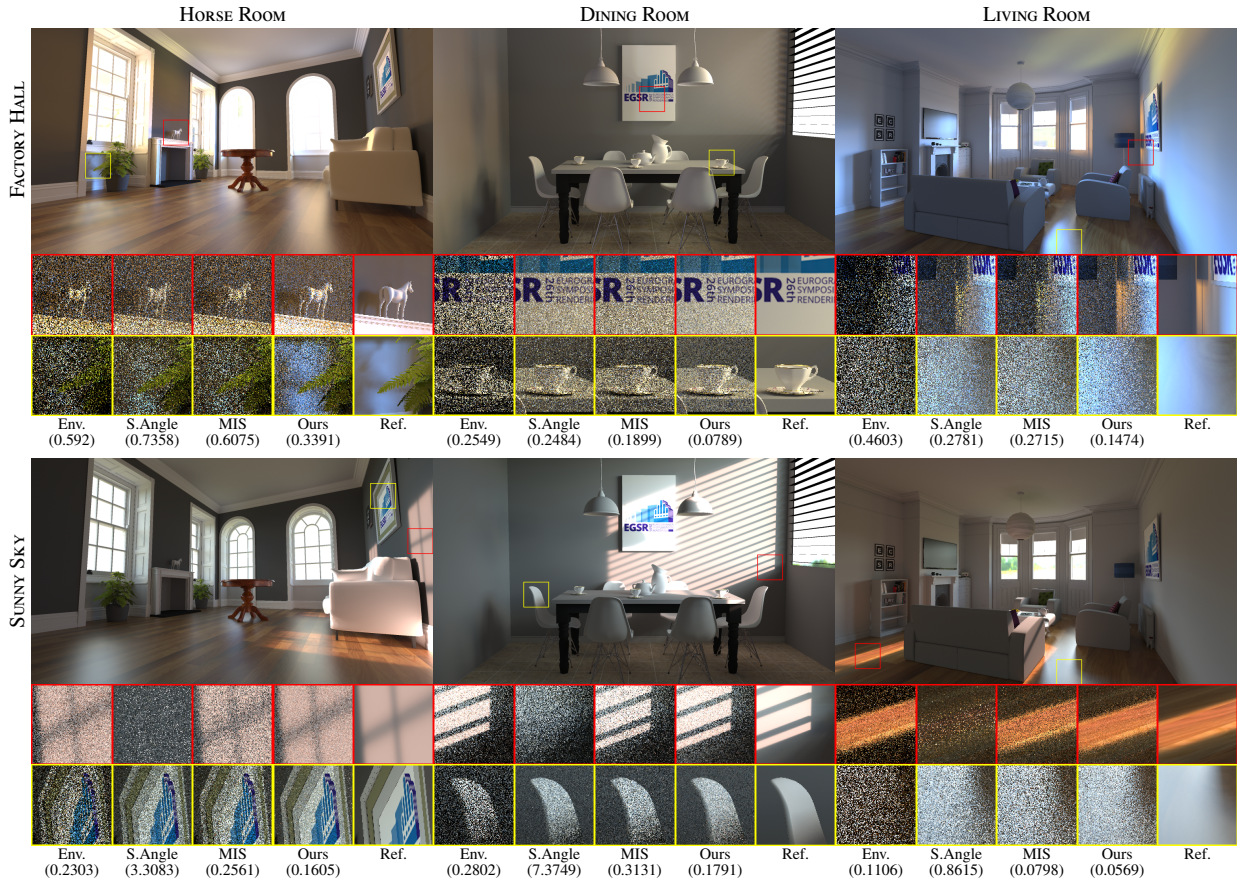


Figure 4: For scenes illuminated by environment maps, our technique provides the best RMSE (in parenthesis) compared to the three baseline methods. We show two zoom-in locations for each technique; please refer to the supplemental material for full-resolution images. Table 1 reports additional performance statistics for all tested scene–environment map combinations.

Distant Lighting. We compare the following baseline approaches for distant light sampling:

- 1) traditional environment map sampling;
- 2) uniform sampling of the solid angle subtended by the portal [UFK13], with portals selected proportional to their solid angle;
- 3) techniques 1) and 2) combined using the one-sample MIS model with balance heuristic [Vea98],

to the following two variants of our algorithm:

- 4) our importance sampling, with portals selected proportional to their solid angle; and,
- 5) our importance sampling, with portals selected proportional to their portal-visible energy.

For each light sampling technique, we perform multiple importance sampling with one additional BRDF sample.

We tested these five approaches in three indoor scenes (HORSE ROOM, DINING ROOM, LIVING ROOM) illuminated by three different environment maps (AMBIENT, FACTORY HALL, and SUNNY SKY). We manually placed one portal behind each

window, fully covering the window opening. The original environment map resolutions were 1500×750 for SUNNY SKY, 1024×512 for FACTORY HALL and 1×1 for AMBIENT. We always evaluate radiance using the full-resolution versions, but re-sample each one to a 512×512 rectified summed-area table for the purposes of importance sampling. In all scenes, the precomputation per portal took 0.1 seconds or less and was insignificant compared to the total render time. Some of the scenes are shown in Figure 4 (please refer to the supplementary material for an exhaustive comparison). The insets depict the noise of each technique with 8 paths per pixel. Note that only our approach can handle all possible scenarios robustly, albeit at the expense of slightly higher sampling cost. We report the *cost* (normalized w.r.t. environment sampling) and various other statistics in Table 1. In order to quantify the different noise characteristics and performance of all sampling techniques in a single representative number, we use the *time-to-unit-variance (TTUV)*—computed as the $MSE \times$ render time—as the primary metric for comparing the different methods. Since all the compared methods are unbiased and use i.i.d. samples, this product represents the time in seconds

required to achieve an MSE value of 1. We report the TTUV ratio w.r.t. the baseline techniques as our *effective speedup*. Please see the supplementary material for images rendered using low-discrepancy sequences.

With portal-visible energy selection, our method is always better than environment map sampling (1.3–8.3× speedup). Uniformly sampling the solid angle of the portal performs very well on the featureless AMBIENT environment map, where we perform roughly the same (0.9× speedup); however, with higher frequency environment lighting (FACTORY HALL and SUNNY SKY) it becomes ineffective, and our method yields 3.3–1486× speedup. The solid-angle sampling is particularly problematic in regions which are directly illuminated (see the red insets for SUNNY SKY in Figure 4) as it can sample the visible sun only by random chance. The MIS baseline technique is somewhat more robust than the other two in isolation, but our approach performs even better across all lighting and scene choices (1.2–4.9× speedup).

IES Rectangular Lights. Importance sampling a rectangular light with a directional emission profile is a dual problem to sampling the environment map visible through a portal. Such light sources are popular since they allow artistic control of both the directional profile and the softness of the shadows (due to the finite area of the light). Applying our importance sampling to direct lighting from such an area light is straightforward: the area light corresponds to the portal and the emission profile to the environment map (sweeping the profile over the light is the same as centering the profile flipped at the shading location). Figure 5 shows the Cornell box with two such light sources where the radially symmetric directional distributions are defined with a 1×64 resolution image in spherical coordinates (resampled to 512×512 in rectified space). We can additionally leverage our probabilistic “portal” selection here to choose which light to sample based on the integrated profile visible through each area light. This greatly reduces noise in regions where the more distant light (w.r.t. the shading point) dominates the illumination due to alignment with peaks of the directional distribution. Using our technique reduces the RMSE of estimating direct illumination by a factor of 3.6. When including global illumination the RMSE still improves by a factor of 1.8.

5. Conclusion, Limitations & Future Work

We presented a practical technique for importance sampling the visible portion of a tabulated directional distribution through a rectangular opening. To make this practical, we rely on a novel rectification of the hemisphere of directions, which always maps the rectangular opening to axis-aligned rectangular regions in the tabulation. Our technique can then be used for two dual rendering problems which are challenging to sample with currently available techniques: importance sampling the environment map visible at a shade point through a portal, or importance sampling the directional emission distribution seen on a quad light from a shade point. Our approach




Figure 5: We compare our method (right side of each image) to solid angle sampling [UFK13] (left side of each image) at equal sample count for the dual problem of importance sampling area lights with directional emission profiles. Our method has only a 22–24% overhead, but produces RMSE values (in parenthesis) that are 1.8–3.6× lower. The false-color visualization on the right shows the allocation of samples to the two light sources using our probabilistic selection based on the portal-visible contribution (top) vs. the baseline’s probabilistic selection based on subtended solid angle (bottom).

is practical and efficient, and results in considerable variance reduction compared to available techniques.

Jacobian and Area Preservation. Since our mapping does not preserve area, we currently account for this by pre-multiplying the resampled environment map by the Jacobian (illustrated in Figure 3). While this is approximately canceled out when dividing by the Jacobian in the solid angle measure PDF during rendering (Equation (12)), some care is needed to avoid numerical instability in the corners where the Jacobian approaches infinity. We did not find this to be a problem in our tests, but it is nonetheless an unfortunate nuisance.

Relation to Ureña et al. Our rectified coordinates share similarities with the parametrization proposed by Ureña et al. [UFK13]. Indeed, our angle α is very similar to their angle ϕ , but, while we parametrize the vertical axis with an analogous angle β , their second coordinate h can be interpreted as the vertical height in cylindrical coordinates. The effective difference between the two approaches is that theirs preserves area, while ours preserves axis-aligned straight lines.

Multi-Product Sampling. Our rectification implicitly samples the product of visibility and lighting by considering only the visible rectangular region. It may be fruitful to consider mapping other terms of the reflection equation—such as the BRDF, cosine term, or tinting texture on the portal—to rectified coordinates and handle the resulting, axis-aligned multi-product e.g. using wavelet importance sampling [CJAMJ05].

Acknowledgements. We thank the reviewers for their helpful suggestions, the artists Jay Hardy and Wayne for the interior scenes in Figure 4 (available at <http://blendswap.com>), and Bernhard Vogl and <http://openfootage.net> for the environment maps used in Figures 1, 3 and 4.

Table 1: We compare three baseline techniques to our method (using two different forms of probabilistic portal selection) on a collection of nine scene–envmap combinations. We report the RMSE, the cost (normalized to envmap sampling), and the TTUV (time to unit variance, i.e. $MSE \times \text{render time}$) in seconds. The best TTUV values are marked in bold. We also list the effective speedup compared to each baseline method.

Envmap	Scene	Envmap Sampling			Portal S.Angle Sampling			Envmap/S.Angle MIS			Ours (solid angle select)				Ours (portal-vis. energy select)			
		RMSE	Cost	TTUV	RMSE	Cost	TTUV	RMSE	Cost	TTUV	RMSE	Cost	TTUV	Eff. Speedup	RMSE	Cost	TTUV	Eff. Speedup
AMBIENT	HORSE	0.3821	1.0	52.87	0.2371	1.5	30.39	0.2948	1.2	39.13	0.2366	1.6	33.02	1.6/0.9/1.2	0.2366	1.6	33.20	1.6/0.9/1.2
	DINING	0.2612	1.0	18.14	0.0998	1.2	3.17	0.1375	1.1	5.47	0.0999	1.3	3.57	5.1/0.9/1.5	0.0999	1.4	3.60	5.0/0.9/1.5
	LIVING	0.2091	1.0	16.45	0.0711	1.5	2.90	0.0940	1.3	4.17	0.0711	1.7	3.16	5.2/0.9/1.3	0.0712	1.7	3.16	5.2/0.9/1.3
FACTORY	HORSE	0.5920	1.0	135.40	0.7358	1.4	293.52	0.6075	1.2	170.40	0.3665	1.5	79.88	1.7/3.7/2.1	0.3391	1.5	68.59	2.0/4.3/2.5
	DINING	0.2549	1.0	18.60	0.2484	1.1	19.68	0.1899	1.1	10.94	0.0789	1.3	2.24	8.3/8.8/4.9	0.0789	1.3	2.25	8.3/8.8/4.9
	LIVING	0.4603	1.0	84.56	0.2781	1.4	44.36	0.2715	1.2	35.92	0.1597	1.6	16.06	5.3/2.8/2.2	0.1474	1.6	13.57	6.2/3.3/2.6
SUNNY	HORSE	0.2303	1.0	20.40	3.3083	1.4	5941.94	0.2561	1.2	30.72	0.2666	1.5	42.15	0.5/141.0/0.7	0.1605	1.6	15.39	1.3/386.1/2.0
	DINING	0.2802	1.0	22.99	7.3749	1.1	17329.34	0.3131	1.0	29.98	0.1791	1.3	11.74	2.0/1476.2/2.6	0.1791	1.2	11.66	2.0/1486.0/2.6
	LIVING	0.1106	1.0	4.84	0.8615	1.5	428.82	0.0798	1.2	3.13	0.0705	1.6	3.13	1.5/137.0/1.0	0.0569	1.6	2.03	2.4/211.1/1.5

References

- [ARBj03] AGARWAL S., RAMAMOORTHY R., BELONGIE S., JENSEN H. W.: Structured importance sampling of environment maps. *ACM Trans. Graph. (Proc. SIGGRAPH)* 22, 3 (July 2003), 605–612. doi:10.1145/882262.882314. 1
- [Arv95] ARVO J.: Stratified sampling of spherical triangles. In *Annual Conference Series (Proc. SIGGRAPH)* (1995), pp. 437–438. doi:10.1145/218380.218500. 2
- [BGH05] BURKE D., GHOSH A., HEIDRICH W.: Bidirectional importance sampling for direct illumination. In *Rendering Techniques (Proc. EG Symposium on Rendering)* (2005), pp. 147–156. doi:10.2312/EGWR/EGSR05/147-156. 1
- [BN76] BLINN J. F., NEWELL M. E.: Texture and reflection in computer generated images. *Commun. ACM* 19, 10 (Oct. 1976), 542–547. doi:10.1145/360349.360353. 1
- [CAM08] CLARBERG P., AKENINE-MÖLLER T.: Exploiting visibility correlation in direct illumination. *Comp. Graph. Forum (Proc. EG Symposium on Rendering)* 27, 4 (2008), 1125–1136. 1
- [CETC06] CLINE D., EGBERT P. K., TALBOT J. F., CARDON D. L.: Two stage importance sampling for direct lighting. In *Rendering Techniques (Proc. EG Symposium on Rendering)* (2006), pp. 103–113. doi:10.2312/EGWR/EGSR06/103-113. 1
- [CJAMj05] CLARBERG P., JAROSZ W., AKENINE-MÖLLER T., JENSEN H. W.: Wavelet importance sampling: Efficiently evaluating products of complex functions. *ACM Trans. Graph. (Proc. SIGGRAPH)* 24, 3 (Aug. 2005), 1166–1175. doi:10.1145/1073204.1073328. 1, 6
- [Cro84] CROW F. C.: Summed-area tables for texture mapping. In *Computer Graphics (Proc. SIGGRAPH)* (1984), pp. 207–212. doi:10.1145/800031.808600. 4
- [Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: Bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Annual Conference Series (Proc. SIGGRAPH)* (1998), pp. 189–198. doi:10.1145/280814.280864. 1
- [GKPS12] GEORGIEV I., KŘIVÁNEK J., POPOV S., SLUSALLEK P.: Importance caching for complex illumination. *Comp. Graph. Forum (Proc. Eurographics)* 31, 2 (2012). doi:10.1111/j.1467-8659.2012.03049.x. 1
- [Gre86] GREENE N.: Environment mapping and other applications of world projections. *IEEE Comput. Graph. Appl.* 6, 11 (Nov. 1986), 21–29. doi:10.1109/MCG.1986.276658. 1
- [Hd14] HEITZ E., D’EON E.: Importance sampling microfacet-based BSDFs using the distribution of visible normals. *Comp. Graph. Forum (Proc. EG Symposium on Rendering)* 33, 4 (2014), 103–112. doi:10.1111/cgf.12417. 1
- [III91] ILLUMINATION ENGINEERING SOCIETY OF NORTH AMERICA: *IES standard file format for electronic transfer of photometric data and related information*, 1991. 4
- [JCI09] JAROSZ W., CARR N. A., JENSEN H. W.: Importance sampling spherical harmonics. *Comp. Graph. Forum (Proc. Eurographics)* 28, 2 (Apr. 2009), 577–586. doi:10.1111/j.1467-8659.2009.01398.x. 1
- [Jen95] JENSEN H. W.: Importance driven path tracing using the photon map. In *Rendering Techniques (Proc. EG Workshop on Rendering)*, 1995, pp. 326–335. doi:10.1007/978-3-7091-9430-0_31. 1
- [KK03] KOLLIG T., KELLER A.: Efficient illumination by high dynamic range images. In *Rendering Techniques (Proc. EG Symposium on Rendering)* (2003), pp. 45–50. 1
- [ODJ04] OSTROMOUKHOV V., DONOHUE C., JODOIN P.-M.: Fast hierarchical importance sampling with blue noise properties. *ACM Trans. Graph. (Proc. SIGGRAPH)* 23, 3 (2004), 488–495. doi:10.1145/1015706.1015750. 1
- [PH10] PHARR M., HUMPHREYS G.: *Physically Based Rendering, Second Edition: From Theory To Implementation*, 2nd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2010. 1, 4
- [RCL*08] ROUSSELLE F., CLARBERG P., LEBLANC L., OSTROMOUKHOV V., POULIN P.: Efficient product sampling using hierarchical thresholding. *The Visual Computer (Proc. CGI)* 24, 7-9 (2008), 465–474. doi:10.1007/s00371-008-0227-y. 1
- [SA07] SUBR K., ARVO J.: Steerable importance sampling. In *Proceedings of the 2007 IEEE Symposium on Interactive Ray Tracing* (2007), IEEE Computer Society, pp. 133–140. doi:10.1109/RT.2007.4342601. 1
- [UFK13] UREÑA C., FAJARDO M., KING A.: An area-preserving parametrization for spherical rectangles. *Comp. Graph. Forum (Proc. EG Symposium on Rendering)* 32, 4 (2013), 59–66. doi:10.1111/cgf.12151. 2, 5, 6
- [Vea98] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, Stanford, CA, USA, 1998. 1, 2, 4, 5
- [VKŠ*14] VORBA J., KARLÍK O., ŠIK M., RITSCHEL T., KŘIVÁNEK J.: On-line learning of parametric mixture models for light transport simulation. *ACM Trans. Graph. (Proc. SIGGRAPH)* 33, 4 (aug 2014). doi:10.1145/2601097.2601203. 1