

Microfacet-based Normal Mapping for Robust Monte Carlo Path Tracing


VINCENT SCHÜSSLER, Karlsruhe Institute of Technology

ERIC HEITZ, Unity Technologies

JOHANNES HANIKA, Karlsruhe Institute of Technology

CARSTEN DACHSBACHER, Karlsruhe Institute of Technology

Classic Normal Mapping
24 seconds



Microfacet-Based Normal Mapping (ours)
27 seconds



Fig. 1. An image rendered at 256spp with classic normal mapping and our microfacet-based normal mapping (default settings). With classic normal mapping, back-facing normals (red) and energy leaks (green) produce artifacts such as black fringes and it can be challenging to obtain the desired appearance without hacking the data or the light transport algorithm. Our new normal-mapping model makes it possible to obtain a detailed reflective sphere not covered by these black regions without breaking light transport.

Normal mapping enhances the amount of visual detail of surfaces by using shading normals that deviate from the geometric normal. However, the resulting surface model is geometrically impossible and normal mapping is thus often considered a fundamentally flawed approach with unavoidable problems for Monte Carlo path tracing, such as asymmetry, back-facing normals, and energy loss arising from this incoherence. These problems are usually sidestepped in real-time renderers, but they cannot be fixed robustly in a path tracer: normal mapping breaks either the appearance (black fringes, energy loss) or the integrator (different forward and backward light transport); in practice, workarounds and tweaked normal maps are often required to hide artifacts.

We present microfacet-based normal mapping, an alternative way of faking geometric details without corrupting the robustness of Monte Carlo path tracing. It takes the same input data as classic normal mapping and works with any input BRDF. Our idea is to construct a geometrically valid microfacet surface made of two facets per shading point: the one given by

the normal map at the shading point and an additional facet that compensates for it such that the average normal of the microsurface equals the geometric normal. We derive the resulting microfacet BRDF and show that it mimics geometric detail in a plausible way, although it does not replicate the appearance of classic normal mapping. However, our *microfacet-based normal mapping* model is well-defined, symmetric, and energy conserving, and thus yields identical results with any path tracing algorithm (forward, backward, or bidirectional).

CCS Concepts: • **Computing methodologies** → **Reflectance modeling**;

Additional Key Words and Phrases: Normal mapping, path tracing, microfacet theory

ACM Reference format:

Vincent Schüssler, Eric Heitz, Johannes Hanika, and Carsten Dachsbacher. 2017. Microfacet-based Normal Mapping for Robust Monte Carlo Path Tracing. *ACM Trans. Graph.* 36, 6, Article 205 (November 2017), 12 pages. <https://doi.org/10.1145/3130800.3130806>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

0730-0301/2017/11-ART205 \$15.00

<https://doi.org/10.1145/3130800.3130806>

1 INTRODUCTION

In this paper, we use the term *normal mapping* to describe the set of techniques that compute the shading of a surface with a perturbed normal instead of the actual geometric normal. The normal is usually provided by a high-resolution texture and allows for increasing the visual complexity of a surface without increasing the complexity of its geometry. A whole body of techniques and representations

originated from the original work on *bump mapping* [Blinn 1978] and they are ubiquitous today in both video games and offline rendering engines. In this paper, we focus on their application for path tracing.

Use of Normal Mapping in Path Tracers. Some offline renderers use tessellation and displacement mapping for geometric details, but normal mapping remains a cheaper alternative that is still widely used in the offline rendering industry, especially for movies. This can be seen in the documentation of path tracers used for major productions, for example Pixar’s Renderman [Pixar 2015], Solid Angle’s Arnold Renderer [SolidAngle 2016], or Guerilla Render [Mercenaries 2016].

Problems of Normal Mapping. Despite its wide adoption in practice, normal mapping remains a “prehistoric” tool in modern frameworks of rendering, inherited from a period where mathematical coherence was less important for rendering. Physically based rendering is obtained by deriving a model of light transport within an appropriate mathematical framework [Veach 1997] which has to form a mathematically consistent whole. The consistency is particularly important for Monte Carlo path tracing where different techniques are often combined together to improve variance reduction. *Symmetry* and *energy conservation* are important constraints for light transport to be well-defined and consistent.

The problem of normal mapping is that it breaks this consistency by definition: it fakes the light transport of a surface by replacing its geometric normal with another perturbed normal. Faking light transport in this way affects both its symmetry and energy conservation. We review the classic problems of normal mapping and their consequences in Section 3.

An Underexplored Problem. These problems are well-known by practitioners who cannot use normal maps in path tracers without making compromises. In order to reduce artifacts, either normal maps are tweaked by artists or workarounds are made in the integrators that prevent the use of advanced rendering algorithms such as bidirectional path tracing. Veach [1996] proposed a way to fix the broken symmetry due to shading normals for bidirectional path tracing. However, the other problems, illustrated for a detailed sphere in Figure 1, remain unsolved, and they are the ones that result in visible artifacts. It is just impossible to obtain a correct or plausible appearance with classic normal mapping and the only way is to replace the normal map by a height map and use tessellation. However, this is not a viable solution: the goal of normal mapping is to have a cheap alternative without geometric or ray tracing overhead.

Position of our Work. To our knowledge, there is no more recent work dedicated to these fundamental problems of normal mapping, although they are well-known and acknowledged to be painful by practitioners. One reason might be the habitual assumption that faking geometric details cannot come without breaking light transport. The objective of this paper is to demonstrate that it is actually possible to model fake geometric details over a flat surface with a correct light-transport formulation. We propose a simple microfacet model to be used as a replacement for a shading normal, and we

derive its masking-shadowing function (Section 4) and its multiple-scattering BRDF with a random-walk approach (Section 5). With this alternative definition of normal mapping, the light transport of the surface is at the same time symmetric (works with bidirectional techniques) and energy conserving (no black fringes nor energy loss). In Section 6 we show that multiple modeling choices are possible and exemplarily explore two of them. In Section 7 we validate our model and propose one configuration with good properties to be used as default.

We call the realization of this idea *Microfacet-Based Normal Mapping* which has the following properties:

- It fixes all the issues of classic normal mapping: asymmetry, black fringes, and violated energy conservation. To our knowledge, it is the first normal mapping model that does not cause problems with Monte Carlo path tracing. As a consequence, it does not replicate the appearance of classic normal mapping, mainly because of lower contrast due to its improved energy conservation. However, our model produces plausible results even in cases where classic normal mapping fails.
- Like classic normal mapping, it is implemented as a simple wrapper that takes an arbitrary normal map and an arbitrary BRDF (potentially with textured parameters) as input.
- Our default microfacet-based normal mapping model has an analytic expression and is about 35% more costly than classic normal mapping in the worst case and has negligible overhead in realistic scenes.
- Our general microfacet-based normal mapping model comes with optional parameters that provide control over the saturation and the contrasts of the normal-map appearance. In the worst case, it is about 70% costlier than classic normal mapping; in realistic scenes, its cost is largely amortized and the overhead is about 10%.
- The general model is evaluated stochastically and thus leads to additional variance in the renders. For some integrators it is also problematic that the model has no analytic sampling probability density function (PDF). Most notably, our model cannot be used with Metropolis Light Transport out-of-the-box. We discuss this issue and possible solutions further in Section 7.3.

2 RELATED WORK

Microfacet Theory. In Section 4 we propose the *tangent facet* microsurface model which is a variant of the V-cavity model introduced by Torrance and Sparrow [1967]. Our derivations are based on the work of Heitz [2014] who reviewed the equations that should be satisfied by microfacet models and their masking-shadowing functions to be geometrically correct. We also build upon the random-walk algorithm by Heitz et al. [2016] that they use for evaluating the multiple-scattering BRDF of surfaces based on the Smith model.

Mathematically Correct Fake Geometric Details. Our idea is conceptually similar to the model of Raymond et al. [2016] for textured scratches: they precompute several scratch-BRDFs and apply them

over the surface. The resulting light transport is well-defined because the average normal of the scratches is exactly the geometric normal, i.e. the geometry of the fake details is consistent with the surface geometry. In contrast, our model uses a normal map to model scratches shaped as simple V-cavities from which we compute a valid BRDF on the fly. Our V-cavity model does not require any precomputation and we can thus use it for the general case where arbitrary normal maps and BRDFs are given as input.

Normal Map LODs. Like other textures, normal maps are usually prefiltered in order to avoid aliasing and reduce the memory footprint when the details are seen from far away. LOD representations for classic normal mapping are based on the same idea: the normal map details are merged into BRDF parameters, typically the “roughness” for glossy materials [Han et al. 2007; Toksvig 2005]. Some techniques use the same idea in other spaces, for instance LEAN Mapping [Dupuy et al. 2013; Olano and Baker 2010] represents the data as a Gaussian in slope space. In this case, the average slope yields a normal, and the covariance of the Gaussian yields the anisotropic roughness parameters of the BRDF. Since our model works with any normal map and any BRDF (and its parameters), it is compatible with all the aforementioned LOD representations.

Glint Rendering. Microfacet-based normal mapping is not relevant for the very specific case of glint rendering, which usually does not suffer from the black fringes and energy conservation problems of classic normal mapping: Glint rendering integrators are required for highly-specular subpixel normal-mapped details [Yan et al. 2014, 2016] and in this case, large normal map footprints are considered for the glints where possibly occurring backfacing normals are virtually outweighed. Note that glint integrators only work for forward path tracing as the calculation of the footprints violates symmetry by definition. As we focus on fixing all problems (including symmetry), our microfacet-based approach is meant as a replacement for cases where classic normal mapping is applied.

3 PROBLEMS CAUSED BY NORMAL MAPPING

Classic normal mapping replaces the geometric normal by a perturbed normal for shading. However, this results in inconsistencies between the actual geometry and the one defined by the shading normals. In particular, this leads to non-symmetric scattering and violates the conservation of energy as described by Veach [1996].

3.1 Non-Symmetry Due to Shading Normals

The use of shading normals corresponds to a non-symmetric modification of the BRDF. This becomes easily visible when using bidirectional methods, e.g. bidirectional path tracing (see Figure 14).

To show this non-symmetry, let us first consider an intersection of an eye subpath with a normal mapped surface shown in Figure 2. For both eye and light subpaths, we refer to the eye direction as ω_i and to the light direction as ω_o . Let f_{ω_s} be the BRDF evaluated with respect to the shading normal ω_s instead of the geometric normal ω_g .

To achieve the desired effect of a shading normal, the foreshortening factor for ω_o must use ω_s instead of ω_g , leading to the modified




Fig. 2. Intersection of an eye subpath. A ray from the incident direction ω_i leaves the surface in the outgoing direction ω_o , which is determined by importance sampling $f_{\omega_s}(\omega_i, \omega_o)\langle\omega_o, \omega_s\rangle$.

BRDF \tilde{f} :

$$\tilde{f}(\omega_i, \omega_o) = f_{\omega_s}(\omega_i, \omega_o) \frac{\langle\omega_o, \omega_s\rangle}{\langle\omega_o, \omega_g\rangle}. \quad (1)$$

This modified BRDF is in general not symmetric, even if the original BRDF f was symmetric.

For bidirectional rendering it would be necessary to efficiently sample the adjoint BRDF, but we consider this to be impractical. When constructing a light subpath, we can usually only importance sample $f_{\omega_s}(\omega_i, \omega_o)\langle\omega_i, \omega_s\rangle$ instead of the modified BRDF to determine ω_i . Consequently, we need to reweight the contribution:

$$\tilde{f}(\omega_i, \omega_o)\langle\omega_i, \omega_g\rangle = f_{\omega_s}(\omega_i, \omega_o)\langle\omega_i, \omega_s\rangle \frac{\langle\omega_o, \omega_s\rangle\langle\omega_i, \omega_g\rangle}{\langle\omega_o, \omega_g\rangle\langle\omega_i, \omega_s\rangle}. \quad (2)$$

When ω_o is nearly perpendicular to ω_g , or ω_i is nearly perpendicular to ω_s , the weighting factor can get arbitrarily large. This causes high variance in renderings and does not solve the problem in practice¹.

3.2 Black Fringes on Surfaces

Shading normals tilt the positive hemisphere of outgoing directions, leading to inconsistencies with the geometric hemisphere, as shown in Figure 3. The contributions of all directions that lie on any of the two negative hemispheres must be discarded. This causes black fringes to appear when the viewing direction is below the hemisphere around the shading normal. It also leads to considerable darkening of the surface, because of the reduced hemisphere of outgoing directions. This is an additional cause of black fringes for highly specular materials when the reflected viewing direction is below the surface.

Veach suggests to extend the BRDF to the entire sphere to eliminate undefined regions of the geometric hemisphere. While this may be reasonable for simple cases like a diffuse BRDF, it is unclear how this extension should be defined for others.



Fig. 3. Problems caused by normal mapping. With a naive approach, light can leak through the surface when evaluating the BRDF with respect to the shading normal (left). Furthermore, the BRDF is undefined for directions below the tilted hemisphere defined by the shading normal (right).

¹ In the Mitsuba renderer, Veach’s correction factor for shading normals was disabled because it comes with unacceptable variance: <https://www.mitsuba-renderer.org/repos/mitsuba/changeset/b766e85cb429>.

3.3 Violation of Energy Conservation

Energy conservation imposes the following constraint on the modified BRDF:

$$1 \geq \int_{\Omega} \tilde{f}(\omega_i, \omega_o)(\omega_i \cdot \omega_g) d\omega_i \quad (3)$$

$$= \frac{\omega_o \cdot \omega_s}{\omega_o \cdot \omega_g} \int_{\Omega} f_{\omega_s}(\omega_i, \omega_o)(\omega_i \cdot \omega_g) d\omega_i$$

We notice that the factor $\frac{\omega_o \cdot \omega_s}{\omega_o \cdot \omega_g}$ can be moved outside the integral. Note that this factor can become arbitrarily large for outgoing directions nearly perpendicular to the geometric surface. Consequently the constraint is not satisfied and the use of shading normals violates the conservation of energy.

3.4 Standard Practices for Preventing Artifacts

We compare our model to two techniques that are commonly used in practice to hide the black fringe artifacts of classic normal mapping. Both methods define the BRDF for incident directions below the positive shading hemisphere:

- The *flipping technique* (Figure 4 left) flips the direction of the shading normal. In this way, incident directions always lie in the positive hemisphere of the shading normal and no incident directions are undefined.
- The *switching technique* (Figure 4 right) switches to an alternative, arbitrarily chosen BRDF. We use a diffuse BRDF with albedo $\rho = 0.5$.

To complement the described techniques, we implement a technique suggested by Keller et al. [2017] that mitigates the effect of sampled directions lying below the geometric hemisphere. It works by changing the shading normal such that the reflection vector of the incident direction is always in the positive geometric hemisphere. While this does not prevent invalid directions from being sampled, it ensures that specular reflections do not go below the surface.

All of these techniques modify the BRDF in a way that further amplifies issues with energy conservation and symmetry of classic normal mapping. Further, sampling the adjoints of these techniques is problematic, because the hemisphere (flipping, [Keller et al. 2017]) or the BRDF (switching) may change depending on ω_i . When tracing from the light, we need to sample ω_i for a given ω_o . Since the hemisphere and BRDF are not fixed before sampling ω_i , choosing a good sampling strategy for ω_i is difficult.




Fig. 4. Standard techniques to prevent undefined directions. The *flipping technique* (left) changes the orientation of the shading normal if it is backfacing from the incident direction. The *switching technique* (right) switches to an alternative BRDF when an undefined direction is encountered.

4 MICROSURFACE MODEL

In order to solve the problems caused by normal mapping, we propose a microfacet model as an alternative to shading normals. Instead of just replacing the geometric normal for lighting calculations, the perturbed normal determines the orientation of microfacets. In this way, the apparent orientation of the surface is changed locally, while maintaining validity of the BRDF through the application of microfacet theory.

4.1 Choice of the Microsurface Profile

We choose a model in which the microsurface is composed of tiny asymmetric cavities consisting of two opposing facets: While the first facet is oriented by the perturbed normal ω_p , the other one is perpendicular to the geometric surface and its corresponding micro-normal $\omega_t := \omega_t(\omega_p)$ is tangent to the geometric surface. Figure 5 shows the resulting microsurface profile, which is akin to the V-cavity model by Torrance and Sparrow [1967] yet with asymmetric cavities.




Fig. 5. Microsurface profile. We add a tangent facet ω_t that compensates for the perturbed normal such that the average normal of the microsurface (ω_p and ω_t weighted by the areas of the respective facets) remains the geometric normal.

We have chosen this configuration because it is the configuration with minimal area that is geometrically correct and produces a look similar to the desired normal map, i.e. where the desired shading normal ω_p has the dominant influence on the appearance:

- It satisfies the geometric constraint that the average normal of the microsurface is the normal of the geometric surface.
- The distribution of microfacet normals is *not* smooth: note that the final appearance of the surface is a result of the microsurface and the input BRDF. Any smooth distribution would act as a low-pass filter on the input BRDF and thus specular appearance would not be possible. If the microsurface remains discrete (Dirac delta distributions) then high frequencies are preserved and the resulting BRDF is similar to the input BRDF.
- This configuration maximizes the surface area with the desired shading normal ω_p and minimizes the remaining area (in this case the surface area with ω_t).
- Lastly, in this configuration the tangent facet ω_t is invisible to half of the incident directions onto the macro-surface, i.e. it does not interact with these directions directly.

4.2 Distribution of Normals

The distribution of normals (NDF) describes the statistical distribution of the microfacets' normals [Heitz 2014]. This is equal to the area of the respective microfacets of a unit surface, which we derive directly from the microsurface profile. Since the model consists

of only two different micro-normals, the distribution of normals becomes:

$$D(\omega_m) = \frac{\delta_{\omega_p}(\omega_m)}{\omega_p \cdot \omega_g} + \frac{\delta_{\omega_t}(\omega_m) \sqrt{1 - (\omega_p \cdot \omega_g)^2}}{\omega_p \cdot \omega_g}. \quad (4)$$

This NDF satisfies the geometric constraint that its projected area onto the geometric normal should equal one [Heitz 2014]:

$$\int_{\Omega} D(\omega_m) |\omega_m \cdot \omega_g| d\omega_m = 1. \quad (5)$$


4.3 Projected Areas and Intersection Probabilities

The projected areas of the facet oriented by ω_p and the tangent facet onto an incident direction ω_i are determined as

$$a_p(\omega_i) = \frac{\langle \omega_i, \omega_p \rangle}{\langle \omega_p, \omega_g \rangle}, \quad (6)$$

$$a_t(\omega_i) = \frac{\langle \omega_i, \omega_t \rangle \sqrt{1 - \langle \omega_p, \omega_g \rangle^2}}{\langle \omega_p, \omega_g \rangle}, \quad (7)$$

and the probabilities that a ray incident from direction ω_i intersects ω_p or ω_t , respectively, are proportional to the projected areas:



$$\lambda_p(\omega_i) = \frac{a_p(\omega_i)}{a_p(\omega_i) + a_t(\omega_i)}, \quad (8)$$

$$\lambda_t(\omega_i) = \frac{a_t(\omega_i)}{a_p(\omega_i) + a_t(\omega_i)}. \quad (9)$$

Lastly, the distribution of visible normals (VNDF) [Heitz 2014] is:

$$D_{\omega_i}(\omega_m) = \lambda_p(\omega_i) \delta_{\omega_p}(\omega_m) + \lambda_t(\omega_i) \delta_{\omega_t}(\omega_m). \quad (10)$$

4.4 The Masking Function

The masking function $G_1(\omega_i, \omega_m)$ determines the fraction of microfacets with normal ω_m that are visible from direction ω_i . Similar to the Cook-Torrance V-cavity masking function [Heitz 2014] our microsurface profile has only two normals and the following cases can occur: (1) either both facets (with normal ω_p and ω_t) are fully visible and the masking function equals 1; or (2) one of them is backfacing ($a_p(\omega_i)$ or $a_t(\omega_i)$ equals 0), and the other one is partially masked (Figure 6).

This enables us to derive $G_1(\omega_i, \omega_m)$: first of all, it has to fulfill the property of conservation of the projected area [Heitz 2014]:

$$\omega_i \cdot \omega_g = \int_{\Omega} G_1(\omega_i, \omega_m) \langle \omega_i, \omega_m \rangle D(\omega_m) d\omega_m, \quad (11)$$

which for our microsurface simplifies to

$$\langle \omega_i, \omega_g \rangle = G_1(\omega_i, \omega_p) a_p(\omega_i) + G_1(\omega_i, \omega_t) a_t(\omega_i). \quad (12)$$

This condition is met by the following masking function which can conveniently be expressed in a single formula:

$$G_1(\omega_i, \omega_m) = H(\langle \omega_i, \omega_m \rangle) \min \left[1, \frac{\langle \omega_i, \omega_g \rangle}{a_p(\omega_i) + a_t(\omega_i)} \right], \quad (13)$$

where $H(-)$ is the Heaviside function.




Fig. 6. The different cases of the masking function. Either both normals are fully visible (left), only ω_p is visible (middle), or only ω_t is visible.

4.5 The Masking-Shadowing Function

For simplicity, we opted for the separable masking-shadowing function given by Equation 14 which assumes no correlation of masking and shadowing:

$$G_2(\omega_i, \omega_o, \omega_m) = G_1(\omega_i, \omega_m) G_1(\omega_o, \omega_m). \quad (14)$$

This choice is motivated by the simple random-walk approach used to compute the multiple-scattering BRDF of the microsurface and is explained in the next section.

5 THE BRDF OF THE MICROSURFACE

In this section, we derive the BRDF of the microsurface given the BRDFs of the perturbed and tangent facets f_p and f_t , respectively.

5.1 The Single-Scattering BRDF

The contribution of facets with a single micro-normal ω_m and BRDF f_m to the BRDF of the entire surface can be obtained from the single-scattering microfacet BRDF [Heitz 2014]:

$$f(\omega_i, \omega_o, \omega_m) = \frac{f_m(\omega_i, \omega_o) \langle \omega_o, \omega_m \rangle \langle \omega_i, \omega_m \rangle G_2(\omega_i, \omega_o, \omega_m) D(\omega_m)}{|\omega_g \cdot \omega_o| |\omega_g \cdot \omega_i|}. \quad (15)$$

The value of the single-scattering BRDF f_1 is the integral of the contribution of each micro-normal over the hemisphere:

$$f_1(\omega_i, \omega_o) = \int_{\Omega} f(\omega_i, \omega_o, \omega_m) d\omega_m. \quad (16)$$

In our case, the cosine-weighted BRDF expression simplifies to

$$f_1(\omega_i, \omega_o) \langle \omega_o, \omega_g \rangle = \lambda_p(\omega_i) f_p(\omega_i, \omega_o) \langle \omega_o, \omega_p \rangle G_1(\omega_o, \omega_p) + \lambda_t(\omega_i) f_t(\omega_i, \omega_o) \langle \omega_o, \omega_t \rangle G_1(\omega_o, \omega_t), \quad (17)$$

which is the sum of the cosine-weighted microfacet-BRDFs multiplied by their respective intersection probabilities with the incident direction and their shadowing function in the outgoing direction. The single-scattering BRDF f_1 is symmetric and energy conserving if f_p and f_t fulfill these properties.

5.2 The Multiple-Scattering BRDF

As shown for the Smith model [Heitz et al. 2016], the multiple-scattering BRDF $f_{\infty}(\omega_i, \omega_o)$ of a microsurface can be defined as the expectation of random-walks on the microsurface. Following this idea, we propose a similar random-walk method illustrated in Figure 7 and presented in Algorithm 1.

Random Walk Algorithm. In order to perform the random walk (and compute its result) we need to track the current ray direction ω_r , its energy throughput e , the total radiance L_o collected from the direction ω_o towards the light source, and the current facet (orientation) ω_m . The random walk starts either at the facet oriented by ω_p by or ω_t , which are chosen according to the respective probabilities $\lambda_p(\omega_i)$ and $\lambda_t(\omega_i)$.

At each intersection, we increase L_o by the radiance reflected by the current facet ω_m weighted by the current ray throughput e and the masking function $G_1(\omega_o, \omega_m)$ in direction ω_o . Then we sample the next direction ω_r from the BRDF of the current facet f_m and multiply the energy throughput by the weight of the sample. The ray leaves the microsurface with a probability given by the masking function $G_1(\omega_r, \omega_m)$, or it intersects the respective other facet.




Fig. 7. Random walk on the microsurface. We sample the first intersection according to probabilities λ_p and λ_t , and the next intersections always occur on the opposite facet. After each intersection, the energy is redistributed uniformly on the facet, consistently with the separable masking-shadowing, and the shadowing function yields the probability of leaving the microsurface.

Algorithm 1 Random walk on the microsurface

```

 $L_o = 0$                                 ▶ radiance collected from  $\omega_o$ 
 $\omega_r \leftarrow -\omega_i$                   ▶ ray direction
 $e \leftarrow 1$                           ▶ ray throughput

 $\omega_m \leftarrow [\mathcal{U} < \lambda_p(\omega_i) ? \omega_p : \omega_t]$  ▶ first intersection
while true do
   $L_o \leftarrow L_o + e f_m(-\omega_r, \omega_o) \langle \omega_o, \omega_m \rangle G_1(\omega_o, \omega_m)$  ▶ eval

   $(\omega'_r, w) \leftarrow \text{sample } f_m(-\omega_r, \omega'_r) \langle \omega'_r, \omega_m \rangle$  ▶ sample
   $\omega_r \leftarrow \omega'_r$                   ▶ update direction
   $e \leftarrow w e$                       ▶ update throughput

  if  $\mathcal{U} < G_1(\omega_r, \omega_m)$  then
    break                                ▶ leave microsurface
  else
     $\omega_m \leftarrow \text{other facet}(\omega_m)$  ▶ intersect other facet
  end if
end while

```

Our algorithm is conceptually similar to that of Heitz et al. [2016], but it is simpler for two reasons:

- We use a *discrete NDF* opposed to continuous NDFs, such as Beckmann or GGX, which need to be importance sampled

at each intersection with the microsurface. As our microsurface is made of only two facets, we sample the first facet using the intersection probabilities of Equations 8, and the next intersection can only occur on the other facet.

- In order to simplify the computation, we have chosen a *separable masking-shadowing function* (Equation 14).

A random walk by explicitly tracing rays (computing geometric ray-facet intersections) and tracking the heights of the intersection points would be consistent with non-separable masking-shadowing functions that model the correlations due to the heights within the microsurface.

Intuitively, assuming separable masking-shadowing means that whenever a ray intersects a facet, its position is uniformly redistributed on the facet surface independently of the height of the intersection point [Heitz 2014]. Due to this assumption we can use the average intersection probability given by the shadowing function $G_1(\omega_o, \omega_m)$ instead of computing explicit ray-facet intersections.

This assumption yields a simple and mathematically valid algorithm as it does not break the symmetry and the conservation of energy of light transport on the microsurface. Furthermore, in the case where there is either no masking or no shadowing (the first case in Figure 6) then the separable and non-separable masking-shadowing functions are equivalent and our random-walk algorithm is equivalent to an explicit ray tracing simulation with geometric ray-facet intersections.

Properties of the Algorithm. In practice, we use Algorithm 1 to either evaluate or sample the cosine-weighted multiple-scattering BRDF $f_\infty(\omega_i, \omega_o) \langle \omega_g, \omega_o \rangle$. Indeed, at the end of the algorithm, the value of L_o yields an unbiased estimate of the cosine-weighted multiple-scattering BRDF:

$$f_\infty(\omega_i, \omega_o) \langle \omega_g, \omega_o \rangle = E[L_o], \quad (18)$$

and the direction ω_r is a sample of it with weight e .

Properties of the Multiple-Scattering BRDF. The BRDF f_∞ defined as the expectation of random walks on the microsurface in (Equation 18) is symmetric and energy conserving:

$$f_\infty(\omega_i, \omega_o) = f_\infty(\omega_o, \omega_i), \quad (19)$$

$$\int_{\Omega} f_\infty(\omega_i, \omega_o) \cos \theta_o d\omega_o \leq 1. \quad (20)$$

Furthermore, it achieves 100% energy conservation:

$$\int_{\Omega} f_\infty(\omega_i, \omega_o) \cos \theta_o d\omega_o = 1, \quad (21)$$

if so do f_p and f_t .

6 CHOOSING THE BRDF OF THE TANGENT FACET

In Section 5 we derived the single- and multiple-scattering BRDFs of the microsurface which depend on the BRDFs f_p and f_t of the facets. Since our goal is to obtain the appearance of a normal map defined by an input perturbed normal ω_p and its BRDF f_p , it is meaningful to keep f_p as provided by the input. However, the choice of the BRDF f_t of the tangent facet is more arbitrary. Ideally, f_t would interfere as little as possible with the average material appearance.

Along this idea, we explore two meaningful options for f_t : using either the same BRDF f_p , or using a purely specular material.

6.1 Choosing a Tangent Facet of the Same Material

The first option is to use the same BRDF for the tangent facet and the perturbed facet, i.e. $f_p = f_t$ (Figure 8). The reasoning is that the tangent facet then contributes to the transport similarly as the surrounding perturbed facets would (assuming that material properties vary slowly across the surface).




Fig. 8. Choosing a tangent facet of the same material.

Evaluation and Sampling. In this case, the evaluation and sampling of the multiple-scattering exactly follow Algorithm 1 with $f_m = f_p$ at each intersection.

Analytic Diffuse Multiple Scattering. If the BRDF is diffuse, i.e. $f_m(\omega_i, \omega_o) = \frac{\rho}{\pi}$, we can compute the multiple-scattering BRDF analytically using a radiosity method. Recall that classic radiosity [Goral et al. 1984] solves the light transport between diffuse patches with uniform energy distribution which is precisely the assumption of the separable masking-shadowing function (discussed in Section 5.2). By inversion of the resulting 2×2 transport matrix (for the perturbed and the tangent facet) we obtain the radiosity of each microfacet after an infinite number of reflections as a function of their initial radiosities, i.e. the direct illumination. By this the multiple-scattering BRDF can be evaluated analytically and we use Algorithm 1 only for importance sampling. We provide more details in our supplemental material.

6.2 Choosing a Specular Tangent Facet

The second option is to use a purely specular BRDF for the tangent facets with

$$f_t(\omega_i, \omega_o) \langle \omega_t, \omega_o \rangle = \delta_{\omega_r}(\omega_o), \quad (22)$$

where $\omega_r = \text{reflect}(\omega_i, \omega_t)$, i.e. ω_r is the reflection of ω_i at the tangent facet with normal ω_t . Using this option, the tangent facet is a reflector of the perturbed facet and the appearance of the resulting microfacet surface is driven by the BRDF f_p and the orientation of ω_p . Note that this choice for f_t also affects the random walk in Algorithm 1 as there is no need to sample the specular reflection.




Fig. 9. Choosing a purely specular BRDF for the tangent facet.

Evaluation and Sampling. Directly applying Algorithm 1 does not work because of the evaluation of the specular BRDF f_t . However, we can easily modify the algorithm for this case. This is illustrated in Figure 10: we evaluate f_p with the direction ω'_o , ω_o reflected at the tangent facet ω_t . We weight this contribution by the intersection probabilities for this path:

- the ray leaving ω_p in direction ω'_o intersects the tangent facet ω_t with a probability given by the complement of the masking function: $1 - G_1(\omega'_o, \omega_p)$, and
- the ray leaves ω_t in direction ω_o without further intersection with a probability given by the masking function $G_1(\omega_o, \omega_t)$.

These modifications of Algorithm 1 are highlighted in blue in Algorithm 2.




Fig. 10. Evaluation of the BRDF in our random-walk algorithm with a specular tangent facet. Since the specular BRDF f_t cannot be evaluated, we replace it by the evaluation of the other BRDF f_p on the reflected outgoing direction ω'_o .

Algorithm 2 Random Walk on the Microsurface with Specular ω_t

```

 $L_o = 0$                                 ▶ radiance collected from  $\omega_o$ 
 $\omega_r \leftarrow -\omega_i$                     ▶ ray direction
 $e \leftarrow 1$                             ▶ ray throughput

 $\omega_m \leftarrow [\mathcal{U} < \lambda_p(\omega_i) \ ? \ \omega_p : \omega_t]$     ▶ first intersection
while true do
  if  $\omega_m = \omega_p$  then
     $L_o \leftarrow L_o + e f_p(-\omega_r, \omega_o) \langle \omega_o, \omega_p \rangle G_1(\omega_o, \omega_p)$     ▶ eval
     $L_o \leftarrow L_o +$                                                             ▶ eval
     $e f_p(-\omega_r, \omega'_o) \langle \omega'_o, \omega_p \rangle (1 - G_1(\omega'_o, \omega_p)) G_1(\omega_o, \omega_t)$ 
     $(\omega'_r, w) \leftarrow \text{sample } f_p(-\omega_r, \omega'_r) \langle \omega'_r, \omega_p \rangle$     ▶ sample
     $\omega_r \leftarrow \omega'_r$                                                             ▶ update direction
     $e \leftarrow w e$                                                             ▶ update throughput
  else
     $\omega_r \leftarrow \text{reflect}(-\omega_r, \omega_t)$     ▶ reflect
  end if

  if  $\mathcal{U} < G_1(\omega_r, \omega_m)$  then
    break                                ▶ leave microsurface
  else
     $\omega_m \leftarrow \text{other facet}(\omega_m)$     ▶ intersect other facet
  end if
end while

```

Analytic Single and Double Scattering. Thanks to the simplifications due to the specular material f_t explained above, the single- and double-scattering BRDF has the analytic expression

$$\begin{aligned} f_2(\omega_i, \omega_o) \langle \omega_o, \omega_g \rangle = & \\ & \lambda_p(\omega_i) f_p(\omega_i, \omega_o, \omega_p) \langle \omega_o, \omega_p \rangle G_1(\omega_o, \omega_p) \\ & + \lambda_p(\omega_i) f_p(\omega_i, \omega'_o, \omega_p) \langle \omega'_o, \omega_p \rangle (1 - G_1(\omega'_o, \omega_p)) G_1(\omega_o, \omega_t) \\ & + \lambda_t(\omega_i) f_p(\omega'_i, \omega_o, \omega_p) \langle \omega_o, \omega_p \rangle G_1(\omega_o, \omega_p), \end{aligned} \quad (23)$$

where

- the first line accounts for the single scattering that contains only the path $i \rightarrow p \rightarrow o$ because the path $i \rightarrow t \rightarrow o$ evaluates to 0 since $\text{reflect}(\omega_o, \omega_t)$ is in the lower hemisphere,
- the second line accounts for the paths $i \rightarrow p \rightarrow t \rightarrow o$,
- and the third line accounts for the paths $i \rightarrow t \rightarrow p \rightarrow o$. Since the reflected incident direction $\omega'_i = \text{reflect}(\omega_i, \omega_t)$ is in the lower hemisphere we have an implicit multiplication by $1 - G_1(\omega'_i, \omega_t) = 1$.

7 RESULTS

7.1 Validation

We rendered all the results using the Mitsuba renderer [Jakob 2010] with the classic normal mapping-plugin **normalmap** provided in Mitsuba as well as with our microfacet-based normal mapping-plugin **normalmap_microfacet** which is provided in our supplemental material. The supplemental document also contains more results.

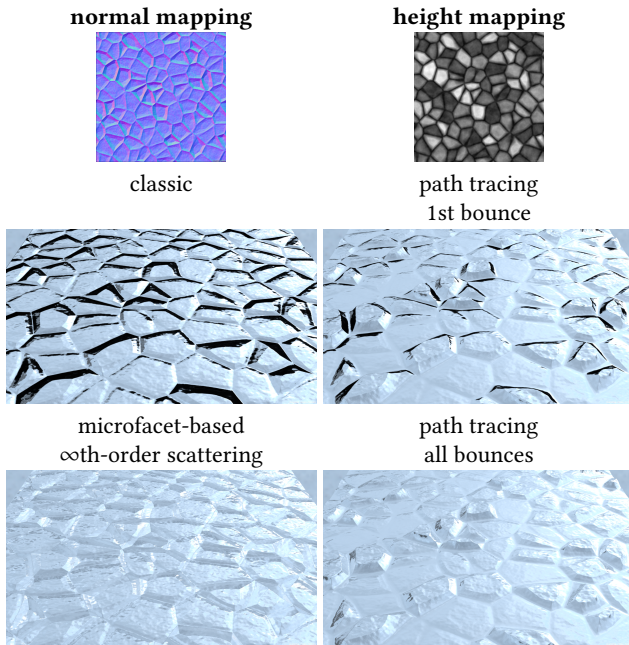


Fig. 11. Comparison of classic normal mapping and single-bounce path tracing to multiple scattering in our model and path tracing.

Removing Artifacts. In Figure 11, we rendered a detailed surface with a specular material with both height and normal mapping.

Classic normal mapping produces black regions where specular reflections are oriented towards the surface due to the problem illustrated in Figure 1 center and Figure 13 (green regions). These black regions are also present on the height-mapped surface if path tracing is cut after the first bounce and they disappear if all the bounces are computed. Thanks to its multiple-scattering component, microfacet-based normal mapping is able to produce a plausible appearance without black regions similar to the full path-traced simulation on the height-mapped surface. This is impossible to achieve with classic normal mapping.

However, note that the microscopic multiple-scattering component of our microfacet-based normal mapping model does not match the macroscopic multiple scattering that occurs on the displaced surface. These are fundamentally different effects and scales, and produce different results. We believe, however, that a multiple-scattering component is mandatory for normal mapping to achieve a *plausible* appearance without artifacts. Figure 13, 16 and 17 show that including up to the 2nd-order scattering already removes almost all artifacts.

Energy Conservation. Figure 13 shows a white furnace test in order to verify that our model is energy conserving. We used a specular BRDF that is 100% energy conserving, and lit the surface with a constant white environment. As explained in Section 3.3, classic normal mapping does not preserve energy resulting in visible darkening. Our BRDF loses less energy thanks to its multiple-scattering term, and accounts for 100% of this energy and passes the white furnace test if all the scattering orders are computed.

Symmetry. In practice Veach’s correction factor [Veach 1996] is often not used in rendering engines because it results in unacceptable variance, however, the consequence of this is non-symmetric light transport. It is also disabled in Mitsuba (see footnote of Section 3.1) and the differences between forward (plugin **path**) and backward path tracing (plugin **ptracer**) easily become apparent (see Figure 14). That is, classic normal mapping does not converge to the same result and thus cannot be used with bidirectional pathtracing techniques as the renderer computes a different result whenever the sampling parameters change. In contrast, our microfacet-based normal mapping is symmetric and both techniques converge to the same result. This problem can also be seen in Figure 15 where we compare with advanced light transport algorithms.

Performance. Our microfacet-based normal mapping can be up to 70% more costly than classic normal mapping with our implementation (Figure 16). However, this is the worst case due to the minimal geometry in this scene with only two triangles. The analytic 2nd-order scattering for the specular tangent facet-model, for example, is 35% slower than classic normal mapping in this case. Figure 17 shows a scene with more complex geometry consisting of 20000 triangles. In this case, the overhead of our normal mapping is already largely amortized and the rendering time is about 10% more than with classic normal mapping.

7.2 Choosing a Default Model

Exploring Appearance. In Figure 16 we compare the appearance obtained with our model for different input BRDFs; please also see

the additional results in our supplemental material. We can observe that the variant of our model using the same material for the tangent facet (Section 6.1) often results in more saturated appearance than the variant using a specular tangent facet (Section 6.2) (best visible for the gold material, top row). This is expected since the absorption factor of the input BRDF (typically Fresnel for microfacet BRDFs or the albedo for diffuse BRDFs) is multiplied twice as often during the random walk if the tangent facet has the same material as the input BRDF. We also note that limiting the scattering order results in some energy loss and produces more contrasted material (of course without the problems of classic normal mapping).

The Best Default Model. In our experiments, the combination that removes most of the artifacts and produces results close to classic normal mapping is the specular tangent facet (Section 6.2) with 2nd-order scattering. Note that this model can be evaluated analytically (Equation 23) and is thus faster than the random-walk approach and comes without additional variance. We believe that these properties qualify it as the default substitution for classic normal mapping. We provide it in our supplemental material as the Mitsuba plugin **normalmap_microfacet_default**.

7.3 Implementation in a Path Tracer

Integration. We integrated our microfacet-based normal mapping in the Mitsuba renderer in the same way as classic normal mapping: it is a simple wrapper over the input BRDF and normal map-texture. It is compatible with any BRDF model (parametric, measured, etc.) and any normal map or equivalent (bump, derivative, etc.).

Default Model. Our analytic default model (Section 7.2, Equation 23, plugin **normalmap_microfacet_default**) works straightforwardly and consistently with any light transport algorithms. This is shown in Figure 15.

General Random-Walk Model. Our general microfacet-based normal mapping (Algorithm 1, plugin **normalmap_microfacet**) with stochastic evaluation does not have a closed-form PDF for the directions sampled by the random walks. However, not knowing the PDF is not a problem for common light transport algorithms such as forward path tracing or bidirectional path tracing.

For computing the importance sampling weights we do not need to evaluate a PDF. Instead, we use a random walk for importance sampling and the BRDF is precisely the expectation of this random walk. By construction of the random walk, the PDF of the generated samples is the cosine-weighted BRDF of a non-absorptive material

$$\text{PDF} = \text{BRDF} \cdot \cos$$

and what remains in the sample weight

$$\text{weight} = \frac{e \cdot \text{BRDF} \cdot \cos}{\text{PDF}} = \frac{e \cdot \text{BRDF} \cdot \cos}{\text{BRDF} \cdot \cos} = e$$

is the energy throughput e that is multiplied by the reflectance after each bounce on the microsurface. Because the BRDF and PDF cancel out in the expression of the importance sampling weight, we do not have to evaluate a PDF in practice to compute these weights. The weight being given by the remaining energy throughput is a classic property of random walk sampling algorithms.

Since our evaluation yields an unbiased stochastic estimate and our importance sampling technique is also unbiased as explained above, we can use *Multiple Importance Sampling* (MIS) [Veach and Guibas 1995] to combine both techniques in a way that remains unbiased. Note that the PDF is usually used for computing the MIS weights of different sampling strategies, but using an approximate PDF does not introduce bias. Indeed, the necessary and sufficient conditions for MIS to be unbiased are:

- the weights of the different techniques should be deterministic,
- the weights should be non-negative, and
- the weights should add up to exactly 1.

Hence, there is no need for an MIS weight to use the exact PDF. For example, common implementations of MIS actually use PDF^2 instead of the PDF itself, which is known as the *power heuristic*. The choice of weights impacts only the variance reduction, but does not introduce bias as long as the three aforementioned conditions are satisfied. Hence, we proceed similarly to Heitz et al. [2016] who approximate the PDF of their multiple-scattering BRDF by the closed-form single-scattering BRDF with a small diffuse contribution to account for multiple scattering. We found this to work well in terms of variance reduction and it makes the model work consistently with forward and bidirectional path tracing, as shown in Figure 15.

It is also possible to make such random-walk BRDFs work with *Metropolis Light Transport* by incorporating the space of the random walks in the path space explored by the algorithms. The contribution and PDF of the paths (random walks) inside the microsurface can be evaluated and mutated by changing the random numbers used to generate them. We explain this in more detail in a supplemental document and leave a further exploration thereof for future work.

Note that our stochastic random-walk evaluation produces results with higher variance than classic normal mapping. Figure 12 shows comparisons of the noise with the same number of samples per pixels. However, note that a non-negligible portion of the additional variance is due to the increased reflectance of the surface and not solely a consequence of the stochastic evaluation.




Fig. 12. Comparison of the variance due to the stochastic evaluation with our general random-walk model.




Fig. 13. Validation of the energy conservation of light transport with a white furnace test. *The material is specular and lit by a constant white environment.*




Fig. 14. Validation of the symmetry of light transport. *We compare forward and backward path tracing; the bottom row shows 4× difference images. The material is a copper conductor with GGX roughness $\alpha = 0.3$. Unlike classic normal mapping, microfacet-based normal mapping produces symmetric BRDFs and we obtain black difference images (the small amount of noise is due to the remaining variance of the renders).*




Fig. 15. Comparison of normal-mapping consistency with different light transport algorithms. *Classic normal mapping does not produce the same image with forward path tracing, bidirectional path tracing, and Metropolis Light Transport. The difference increases considerably when standard practices for fixing artifacts are used. Our default microfacet-based normal mapping, which is an analytic model, works consistently with any light transport algorithms. Our general microfacet-based normal mapping with stochastic evaluation works straightforwardly with forward and bidirectional path tracing. Making it work with MLT is possible in theory, but requires modifications to the integrator. Note that the small amount of noise is due to the remaining variance of the renders.*















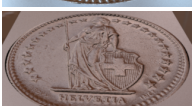
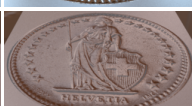












classic			microfacet-based			
classic	switch (3.4)	flip (3.4)	same material tangent facet (6.1)		specular tangent facet (6.2)	
			2nd-order	∞ th-order	2nd-order	∞ th-order
29s	30s	30s	45s	49s	44s using Algo. 2 39s using Eq. (23)	47s
						
						
						
						

Fig. 16. Appearances and performance of our model. We compare the appearances achieved by our model with a tangent facet made of the same material (4th and 5th column, Section 6.1) or specular material (6th and 7th column, Section 6.2). The timings are measured for the gold material shown in the first row, but they are roughly equal across materials shown. We measured the 2nd-order scattering of the specular facet-model using both the random-walk evaluation and its analytic expression.


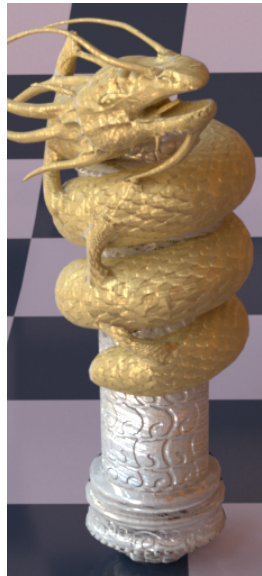

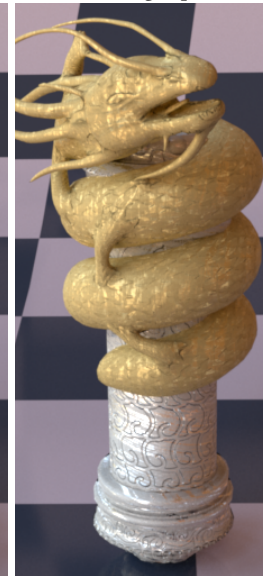
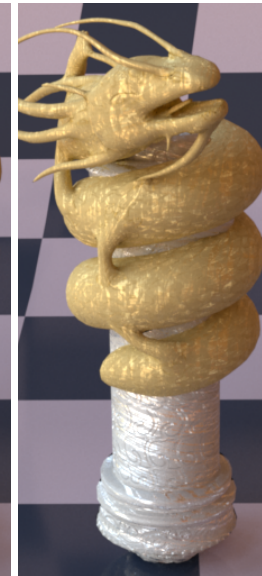
classic			microfacet-based specular tangent facet (6.2)	
classic	switch (3.4)	flip (3.4)	2nd-order scattering	∞ th-order scattering
190s	194s	192s	203s using Algo. 2 202s using Eq. (23)	211s
				

Fig. 17. Our model applied on a more complex geometric model (20000 triangles) with textured GGX roughness. Classic normal mapping (left) shows distracting black fringing. Note the little overhead of our microfacet-based normal mapping in geometrically more complex scenes.

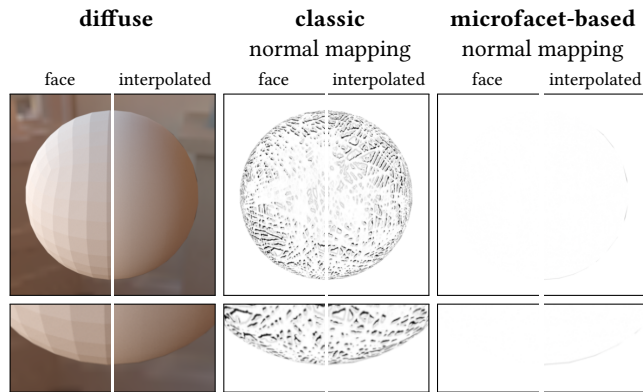


Fig. 18. The problem with interpolated vertex normals. The images show a sphere mesh with face normals or interpolated vertex normals. The center images show a furnace test with classic normal mapping and an energy-conservative BRDF. As expected, a lot of energy is lost due to normal mapping. Our normal-mapping model with infinite scattering order achieves 100% energy conservation but also fails the white furnace test if vertex normals are interpolated (please zoom in, bottom right). This is a separate problem not addressed in this work.

8 CONCLUSION

In this paper, we have shown that faking geometric details on flat surfaces without breaking light transport is actually possible and proposed a new definition of normal mapping based on insights from microfacet theory. Our microfacet-based normal mapping can be used in Monte Carlo path tracing without compromises. It is not subject to energy loss and asymmetry and enables a variety of appearances that were not possible with classic normal mapping.

We demonstrated that microfacet-based normal mapping is practical and is able to fix the artifacts of normal mapping. We believe that it could be easily adopted by practitioners and make their lives easier, especially in production rendering for movies where a lot of time is allocated to tweaking problematic data and hiding artifacts. We hope that despite the changed appearance and stochastic evaluation, our model will prove to be usable as a replacement for classic normal mapping.

Modern Monte Carlo rendering frameworks suffer from other inconsistencies. For instance, interpolating the vertex normals of a mesh is common to make it appear smooth. However, interpolated normals also fake the light transport and result in energy loss and other artifacts such as the *shadow-ray terminator* [Woo et al. 1996]. The case shown in Figure 18 is a white furnace test with a diffuse material and a white environment. We can see that our normal mapping-model does not solve the problems related to interpolated vertex normals – the problems they introduce are different and remain yet to be solved. However, we have shown that the construction of geometrically correct models, such as microfacet surfaces, can be a fruitful approach to this category of problems.

ACKNOWLEDGEMENTS

We thank Jonas Wagner (<https://29a.ch>) for providing the normal map of the coin.

REFERENCES

- James F. Blinn. 1978. Simulation of Wrinkled Surfaces. *Computer Graphics (Proc. SIGGRAPH)* 12, 3 (Aug. 1978), 286–292.
- Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, Fabrice Neyret, and Victor Ostromoukhov. 2013. Linear Efficient Antialiased Displacement and Reflectance Mapping. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 32, 6 (Nov. 2013), 211:1–211:11.
- Cindy M. Goral, Kenneth E. Torrance, Donald P. Greenberg, and Bennett Battaile. 1984. Modeling the Interaction of Light Between Diffuse Surfaces. *Computer Graphics (Proc. SIGGRAPH)* 18, 3 (Jan. 1984), 213–222.
- Charles Han, Bo Sun, Ravi Ramamoorthi, and Eitan Grinspun. 2007. Frequency Domain Normal Map Filtering. *ACM Trans. Graph. (Proc. SIGGRAPH)* 26, 3 (2007), 28:1–28:12.
- Eric Heitz. 2014. Understanding the masking-shadowing function in microfacet-based BRDFs. *Journal of Computer Graphics Techniques* 3, 2 (2014), 32–91.
- Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016. Multiple-scattering Microfacet BSDFs with the Smith Model. *ACM Trans. Graph. (Proc. SIGGRAPH)* 35, 4 (July 2016), 58:1–58:14.
- Wenzel Jakob. 2010. Mitsuba renderer. (2010). <http://www.mitsuba-renderer.org>.
- Alexander Keller, Carsten Wächter, Matthias Raab, Daniel Seibert, Dietger van Antwerpen, Johann Korndörfer, and Lutz Kettner. 2017. The Iray Light Transport Simulation and Rendering System. *arXiv preprint arXiv:1705.01263* (2017).
- Mercenaries. 2016. *Guerilla Render 1.4.2 documentation*. http://guerillarender.com/doc/1.4/Library_Attributes_NormalMap.html.
- Marc Olano and Dan Baker. 2010. LEAN Mapping. In *Proc. ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. 181–188.
- Pixar. 2015. *RenderMan 20 Documentation*. <https://renderman.pixar.com/resources/current/RenderMan/PxrNormalMap.html>.
- Boris Raymond, Gaël Guennebaud, and Pascal Barla. 2016. Multi-scale Rendering of Scratched Materials Using a Structured SV-BRDF Model. *ACM Trans. Graph. (Proc. SIGGRAPH)* 35, 4 (July 2016), 57:1–57:11.
- SolidAngle. 2016. *Arnold for Maya User Guide: Bump Mapping*. <https://support.solidangle.com/display/AFMUG/Bump+Mapping>.
- Michael Toksvig. 2005. Mipmapping Normal Maps. *Journal of Graphics Tools* 10, 3 (2005).
- K. E. Torrance and E. M. Sparrow. 1967. Theory for Off-Specular Reflection From Roughened Surfaces. *Journal of the Optical Society of America* 57, 9 (Sept. 1967), 1105–1114.
- Eric Veach. 1996. Non-symmetric Scattering in Light Transport Algorithms. In *Proc. Eurographics Workshop on Rendering Techniques*. 81–90.
- Eric Veach. 1997. *Robust Monte Carlo methods for light transport simulation*. Ph.D. Dissertation. Stanford University.
- Eric Veach and Leonidas J. Guibas. 1995. Optimally Combining Sampling Techniques for Monte Carlo Rendering. In *Proc. SIGGRAPH*. 419–428.
- Andrew Woo, Andrew Pearce, and Marc Ouellette. 1996. It's Really Not a Rendering Bug, You See. *IEEE Computer Graphics and Applications* 16, 5 (Sept. 1996), 21–25.
- Ling-Qi Yan, Miloš Hašan, Wenzel Jakob, Jason Lawrence, Steve Marschner, and Ravi Ramamoorthi. 2014. Rendering Glints on High-resolution Normal-mapped Specular Surfaces. *ACM Trans. Graph. (Proc. SIGGRAPH)* 33, 4 (July 2014), 116:1–116:9.
- Ling-Qi Yan, Miloš Hašan, Steve Marschner, and Ravi Ramamoorthi. 2016. Position-normal Distributions for Efficient Rendering of Specular Microstructure. *ACM Trans. Graph. (Proc. SIGGRAPH)* 35, 4 (July 2016), 56:1–56:9.