EDITOR: Jim Foley, foley@cc.gatech.edu

DEPARTMENT: DISSERTATION IMPACT

Realistic Rendering in "Details"

Ling-Qi Yan, *University of California, Santa Barbara, CA, 93106, USA*

FROM THE EDITOR

Lingqi Yan received the 2019 SIGGRAPH Outstanding Doctoral Dissertation Award.

Rendering is far from solved. Even today, the rendered results still look artificial and overly perfect. To make rendering more realistic, we need details. However, rendering a complex surface with lots of details is far from easy. Traditionally, the surface microstructure is approximated using a smooth statistical distribution, but this ignores all the details on the surface, completely eliminating the "glinty" visual effects that are easily observable in the real world. While modeling the actual surface microstructure is possible, the resulting rendering problem is prohibitively expensive using Monte Carlo point sampling. We consider the highly complicated distribution of normals on a surface patch seen through a single pixel, and evaluate this actual distribution efficiently with closed-form solutions, in both geometric and wave optics. Results show complicated, temporally varying glints from materials such as bumpy plastics, brushed and scratched metals, metallic paint and ocean waves—bringing the interesting and important details to Computer Graphics for the first time.

For every complex problem, there is an answer that is clear, simple, and wrong."[a] This statement is suitable for most of the current surface reflectance/appearance models for different materials. Traditional rendering techniques represent materials using smooth Bidirectional Reflectance Distribution Functions (BRDFs) describing how light is reflected after interacting with these materials. Since they use smooth BRDFs, these techniques generate perfectly smooth appearances. However, the real world is imperfect. Bumps, flakes, and dents can be seen everywhere. These details introduce variance and are key to the realism of the appearance (Figure 1).

The smooth BRDF concept has been standard for nearly four decades, prior to our work which introduces a more statistical discrete version of the BRDF, and enables rendering of detailed glints from complex surfaces, such as metallic flakes and scratches. These details are either procedurally generated as random processes, or defined using extremely high resolution normal maps specifying the surface normals at different places. Thus, the surfaces are essentially represented using tiny microfacets explicitly.

However, existing microfacet models[1] in Computer Graphics use statistics to represent the distribution of normals, known as normal distribution functions (NDFs). In these models, an NDF is considered as a smooth probability distribution, such as a Gaussian function. Such a smooth NDF results in a smooth appearance, eliminating all the details. In contrast, we compute the actual NDF within a surface patch P covered by each pixel as a P-NDF , introducing an unprecedented level of detail that was never dealt with in Computer Graphics previously. For this reason, we name this line of our research *detailed rendering*.

---

[a] H. L. Mencken.

0272-1716 © 2021 IEEE
Digital Object Identifier 10.1109/MCG.2021.3077918
Date of current version 12 July 2021.

FIGURE 1. Traditional rendering (left) versus detailed rendering from complex surfaces (right). Note how the glints from the tiny brushed dents and scratches dramatically improve realism.

Furthermore, we also notice that the surface details are usually small (in the magnitude of micrometers). In this level of granularity, the light should be treated as waves rather than straight lines. So, in this article, we present two different solutions in geometric optics and wave optics, respectively. Both are able to produce convincing details, while wave optics are further able to correctly generate diffraction effects such as colors from compact discs and dull polished metals.

Defining the Details

To render the details, we first need to define them on surfaces. We refer to texture mapping, a well-known technique to specify per-point properties on a surface. We map high-resolution textures (normal maps) onto surfaces, to define the surface normal of every microfacet. As one might expect, since the microfacets are in micro scale, we need extremely high-resolution normal maps to define every fine detail. For example, the normal map on the snail's shell in Figure 2 has a resolution of 200K × 200K.

Though acquiring the normal maps is orthogonal to our rendering method, it is useful to point out that in practice we usually cannot afford to store such huge normal maps. One simple way to alleviate the storage problem is to create smaller (2K × 2K in our experiments) normal maps that can be seamlessly tiled. To do this, several methods are available. The inverse Fourier transform method[2] is able to generate tileable noise-like bumps, and texture synthesis methods[3] are able to turn measured normal maps into seamlessly tileable patches.

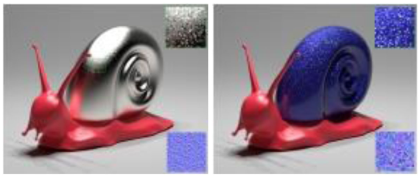

FIGURE 2. Using different kinds of normal maps to define different types of details on a surface. (Left) Isotropic noise. (Right) Metallic flakes.

Despite its high resolution, a normal map is Still, in essence, a texture. So the normals can only be defined per texel, i.e., on a discrete grid with finite resolution. To get continuous normals over an entire surface, we assume that the normal maps are bicubic interpolated. The bicubic interpolation happens on the fly without introducing further storage overhead. That is, whenever any point s is queried on the normal map, a smoothly interpolated normal n(s) will be returned immediately.

According to the microfacet theory, the microfacets do not exhibit any roughness. They are essentially tiny mirrors that are perfectly specular. And it is the distribution of their normals, or NDF, that determines the roughness as well as the full appearance over the entire surface. During rendering, the BRDF is evaluated with pairs of incident-outgoing directions, and the microfacet theory turns each BRDF evaluation into a query on the NDF at the half-vector direction h, midway between the incident and outgoing directions.

In other words, we assume that the BRDFs and NDFs are exchangeable.

DIFFICULT RENDERING PROBLEM

With the definition of microfacets' normal over a surface, we are now ready to render all the details. Before we proceed, we first specify the exact rendering problem that we want to solve. Our focus here is the direct illumination of a complex specular surface, illuminated by a point light source and observed by a pinhole camera.

This is the most difficult configuration, producing the sharpest visual effects of the details. We will deal with other types of light transport later.

The most straightforward approach in rendering is path tracing. We use this standard approach to compute direct illumination on a complex specular surface, tracing a set of rays through each pixel, evaluating the normals of the hit points, and shading the hit points from the light source. Therefore, path tracing is essentially a point sampling approach. However, this classic

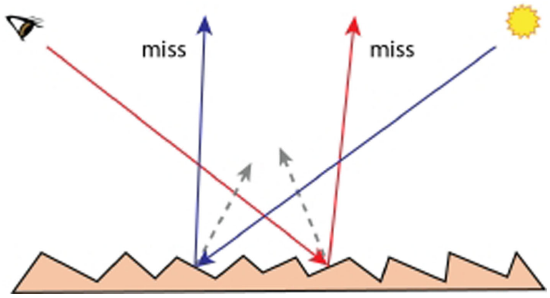

FIGURE 3. Path tracing, either from the camera or from the light source, is not able to find a valid path with perfectly specular microfacets.

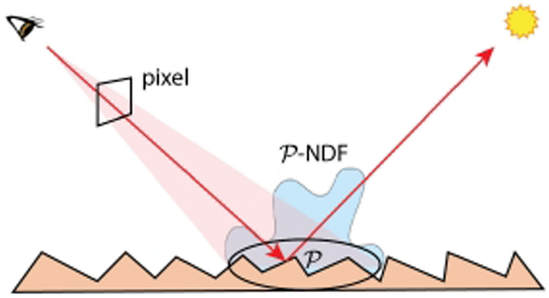

FIGURE 5. We consider the actual P-NDF in a surface patch P as seen through each pixel.

method fails at rendering detailed glints (Figure 3). Similarly, bidirectional approaches[4,5] would also fail. Therefore, multiple importance sampling (MIS)[4] would not help either. This is because MIS requires at least one of the candidate sampling approaches to work well, but there are none in our case.

As we can see, the main problem is the pure specularity that makes point sampling of a complex surface impossible to find valid paths connecting the light and the camera. Consequently, one might think that the problem could be made practical by simply introducing a small amount of roughness, namely the intrinsic roughness, to each micro-facet. However, this is not true. As Figure 4 illustrates, even with the intrinsic roughness, these glinty details are still nearly impossible to render using path tracing in a reasonable amount of time.

One pixel may contain thousands of discrete tiny highlights. To obtain a noise-free image, path tracing needs to find them all. This is a goal too ambitious for point sampling.

PER-PIXEL CONFIGURATION

To surmount the point sampling problem, consider a surface patch P seen through a single pixel and its P-NDF—the actual, unsimplified NDF within this patch P (an example is shown in Figure 5). The P-NDF can then be plugged into any existing microfacet BRDF model, replacing the classic smooth, statistical normal distributions.

The surface patch P covered by a pixel is also known as the pixel's footprint. The shape of a footprint can be irregular, and we simply approximate it with a 2-D Gaussian, with the help of Ray Differentials[6] to determine its size.

Given a pixel's footprint P, its P-NDF can be easily visualized by binning. The P-NDF is usually represented as an image on a projected unit disk. The grid of pixels in this image gives us a natural subdivision of the P-NDF into small directional bins, where each bin records the averaged value of the P-NDF within a small range of directions. The binning method is simple: we repeatedly importance sample a point on the patch P, take its normal, perturb it by the intrinsic roughness, and throw it into a bin according to the perturbed direction.

When we visualize P-NDFs in this way with various sizes of pixel footprints, as Figure 6 shows, we immediately notice that the distributions are far from smooth. In fact, very large patches P are required to make the P-NDFs converge to the smooth functions that we


Our method (2.2 mins) Naive sampling (2 hours) Zoom-in of a single pixel

FIGURE 4. Even with the intrinsic roughness, naive pixel sampling still fails at rendering complex specular surfaces under point lights. The highlights are too small to be efficiently hit by uniform pixel sampling, and are too many to be completely found, which is obvious from the zoomed pixel on the right. (Left) Our method (2.2 min). (Middle) Naive sampling (2 h). (Right) Zoom-in of a single pixel.

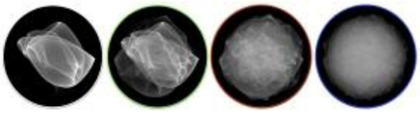

FIGURE 6. Our P-NDFs have sharp features. They converge to smooth distributions as the size of patch P grows, from left to right.

know from microfacet models. Also, as we can tell from Figure 7, different types of normal maps lead to various interesting P-NDF shapes. The only thing they have in common is that they all have sharp features and are certainly not easily approximated by smooth functions. Therefore, any smooth approximation to the true P-NDF[7–10] cannot fully reproduce details.

So long as a P-NDF can be visualized on an image with enough resolution, the binning method can be treated as ground truth. However, P-NDFs generated this way cannot be used directly for rendering. The key problem is that for actual rendering, we want to evaluate the P-NDF at specific places. This is analogous to rendering with any BRDF—we never need to generate the entire 4-D BRDF, but just want to query its value for specific incident and outgoing direction pairs.

For NDFs, as introduced earlier, they are queried at half-vectors h, the mid-way directions between incident and outgoing directions. Clearly, it would be extremely inefficient to use the binning approach here, wasting all but a single bin. Indeed, this is equivalent to what a standard renderer would do, trying to hit a tiny light source by chance.

PRACTICAL RENDERING SOLUTIONS

As analyzed earlier, the fundamental rendering problem is to efficiently query the P-NDF at a given direction s. The answer is such an integral

$$D\,(s) = \int_{\mathrm{R}^2} G_p(u)\; G_r(n(u) - s)\; \mathrm{d}u. \qquad (1)$$

Behind this difficult integral is a very intuitive idea: within a pixel's coverage G_p, we look at every microfacet at position u, and ask if its normal n(u) is close enough to the direction s that we query, considering the intrinsic roughness G_r. In this way, we are able to evaluate the density of a single normal pointing anywhere on the patch. Moreover, the P-NDF is never fully constructed, and is only evaluated for a single direction at a time.

Significant effort is devoted to solve this integral analytically and efficiently. These include triangulating the normal maps,[11] reinterpreting the compound function Gr(n(u)) as Gaussian mixtures,[12] as well as pruning uncontributing microfacets in batches using position-normal hierarchies.[11,12] We refer readers to these publications for more technical details. Figure 8 shows more results.

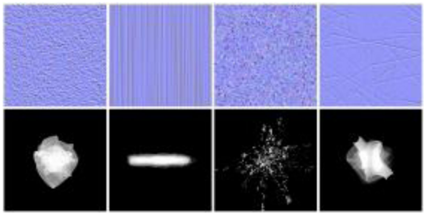

FIGURE 7. Visualization of different P-NDFs for different types of normal maps.

One thing worth noticing is that all these different solutions are aimed at our previously defined goal: direct illumination under point lighting. But the full rendering usually contains other types of light transport, especially multiple bounces of light and illumination under other kinds of lighting conditions. These types of light transport, however, will greatly blur out all the details. With this observation, we simply render multiple bounces of light transport by replacing our detailed surfaces by those defined using microfacet models with perceptually similar roughness. It is even more straightforward to render the direct illumination under other lighting conditions, especially environment lighting: regular path tracing would suffice, because now the mirror reflected rays are able to hit somewhere on area lights and the environment, unlike the case with a point light.

WAVE OPTICS: ONE MORE STEP FORWARD

So far, we have presented practical approaches for accurate and efficient rendering of detailed surfaces, and we are already able to render plausible results. Theoretically, with these methods, we should be able to match real world appearance exactly. However, when we look at the real photos in Figure 9, the observation is quite unexpected.

In these photos, there are colored details even if the objects are illuminated by a white light source. This is an interesting phenomenon that we immediately conclude as impossible for traditional geometric optics—it will produce only white highlights under white light source.

This observation leads to three questions. First, if geometric optics has its limitations at such fine level of detail, is wave optics able to give us more accurate results? Second, can we design a rendering algorithm based on a wave optics model, but still be able to keep spatially varying high-resolution details? Third,

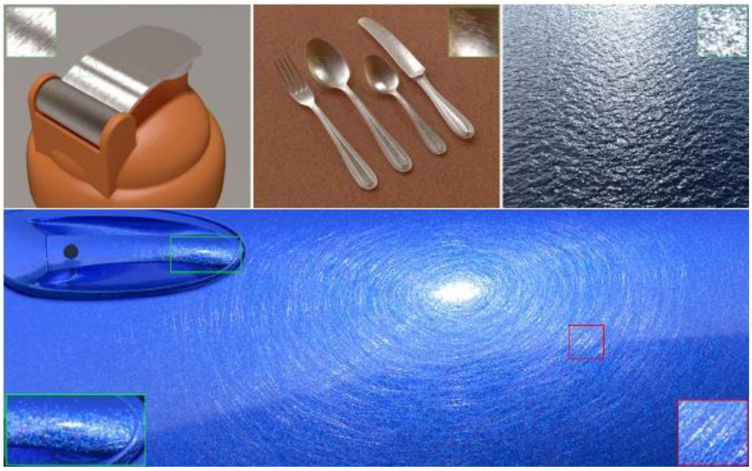

FIGURE 8. More results generated using our practical rendering algorithms. (Top row) using [11]. (Bottom row) using [12].

how close is geometric optics to wave optics, apart from the color difference? Our work is the first to consider all these questions in full generality.

We start by modeling the surfaces as arbitrarily discretized heightfields. This is slightly different from our previous use of normal maps, because now we not only care about the facing directions of each microfacet, but also its relative height with regard to the averaged smooth surface. In other words, a heightfield is able to define the height shift H(s) at every point on a surface. Therefore, when coherent light (traveling "in parallel," starting at the same phase) reaches a small patch P on a surface, different places will introduce different distances that the light travels.

This causes spatially varying phase shifts R(s) for each wavelength λ in reflected waves, which then interfere with each other, introducing colors.

The phase shifts allow us to compute P-BRDFs (the accurate BRDF within a patch P) using wave optics. In wave optics, light is described by complex-valued fields. Scalar diffraction models, such as Harvey-Shack[13] or Kirchhoff,[14] can be used to estimate the reflected field from a rough surface.

FIGURE 9. Real photos of the back of a laptop (left) and an aluminum patch (right), illuminated with only one LED light on a mobile phone. Even though we generally consider these objects white or colorless, the colors introduced by wave optics are always visible.

Complex as the theories are behind the models, the resulting wave BRDF is astonishingly elegant, which is simply related to the Fourier transform of the spatially varying phase shift R(s).

Solving this wave BRDF, however, introduces several new challenges. First, in order for the wave optics to take effect, the microfacet granularity microfacets must be comparable to common wavelengths of the light. This requires an even more detailed definition of a surface as compared to geometric optics. Second, like any other types of BRDFs in rendering, the wave BRDF needs to support point query, i.e., given a specific pair of incident and outgoing directions, we just need to query one value from the wave BRDF rather than the entire BRDF. Third, unlike in geometric optics, all different places within the coherence area, where we assume the light is coherent, will contribute to wave BRDFs. The contributions from different parts of the surface sum nonlinearly (due to interference), strengthening, weakening, and canceling each other to create the characteristic diffraction effects of wave optics. As we can see from Figure 10, apart from being colorful, BRDFs predicted using wave optics all exhibit discontinuous patterns and distinct details.

Thus, when applied in the rendering process, they result in colorful, detailed, and even more realistic results (Figures 11 and 12).

PAST, PRESENT, AND FUTURE

These methods have made promising first steps toward modeling and rendering of visual appearance at real world complexity.

This level of detail was never dealt with in Computer Graphics previously, and our work has brought

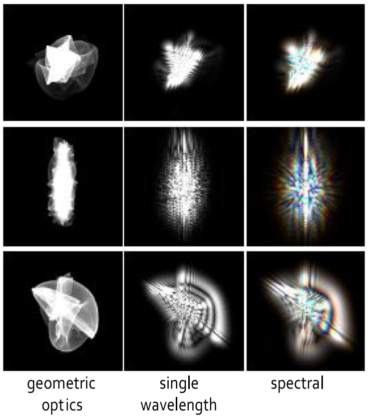

FIGURE 10. Visualizations of the 2-D outgoing BRDF slices with a fixed incident direction perpendicular to the surface. Top: isotropic bumps; middle: brushed; bottom: scratched.

detailed rendering to computer graphics for the first time, bringing out the details and redefining the BRDF concept after nearly four decades of its statistical use. The idea of rendering complex microstructure has also become a rapidly growing subfield[15–18] in rendering. And now, we can conclude that there are no more smooth surfaces in Computer Graphics.

An exciting extension of our work is to provide the ability to represent surface features at all scales with the appropriate type of model: large features can be handled with geometry; smaller ones down to a fraction of a millimeter can be represented using geometric normal maps; and features down to wavelength scale are represented as diffracting height fields. Smaller features than that are not optically relevant and are not needed in any visual simulation. It would also be interesting to bring our approach closer to interactivity with further approximations. An extension to displacement maps would be possible as well.

We could also explore related glinty phenomena caused by refraction, seen, e.g., in snow, hair, waterfalls, fabrics, or plant cellular structures.

However, instead of seeing more problems being solved, we are more enthusiastic to introduce new problems. First of all, an immediate problem to detailed rendering is revealed.

That is, even for now, we still do not know what is ground truth or a trustable reference in wave optics. We make efficient rendering possible only because we base our method on the Harvey-Shack and Kirchoff theories, which make approximations to the general Maxwell's equations. In that sense, only the Maxwell's equations can be considered as the ground truth. However, accurately solving them could be prohibitively slow, even on a tiny patch of surface.

Second, more problems can be brought to other research areas in computer science. For example, the glint patterns rendered from the details can be easily judged by humans as being quite different than noise. However, they may not be easily distinguished from noise by computers, which indicates that they could be a challenge for Computer Vision tasks. For the same reason, producing plausible and clean results may also become a disaster for denoising algorithms when the details exist. And in general, our detailed rendering breaks the assumptions of locally low rank and low frequency. This will bring challenges to those significantly relying on these assumptions, especially Machine Learning based approaches.

However, all these problems and challenges are certainly new opportunities in research, and may stimulate deeper observations and explorations.

Third, our detailed rendering work also negates the claim that Computer Graphics problems will be automatically solved by simply waiting for hardware development. We have proved that tracing rays to/from microstructure is essentially equivalent to point sampling Dirac delta functions blindly, which in theory takes infinite time to converge. Moreover, the growth of hardware's computational power is far from satisfactory. For example, it is general practice that the movie/animation industry would spend thousands of CPU core hours to generate

FIGURE 11. Left and top right: Rendering of a laptop with a point light and environment lighting using our wave optics method. Bottom right: A photograph of a MacBook (around 20 cm × 4 cm region) lit by a small LED light in a dark room. Our method is able to produce a colorful appearance that is perceptually similar to the photograph.

FIGURE 12. Rendering comparison between geometric optics and wave optics on a cutlery scene. Note that even with a single fixed wavelength, wave optics still generates a significantly different appearance as compared to geometric optics.

one frame in a film with 4K resolution (3840 × 2160). Without new technical developments, it will take decades or even centuries before GPUs are able to generate a 2020-film-level image in real time (30 frames per second).

To conclude, even with all our previous efforts, realistic rendering is still far from solved. It is more and more demonstrated that both extremes—realism and speed—have to be satisfied in order to bring people believable virtual contents: it requires even deeper studies of our current rich visual, mathematical, and physical world, but this will probably break many of the assumptions that we make to trade quality for performance. Thus, even more interesting and open problems await us. We sincerely look forward to the day when Computer Generated Imagery (CGI) finally overcomes these difficulties and opens up a view of a new world.

REFERENCES

1. K. Torrance and E. Sparrow, “Theory for off-specular reflection from roughened surfaces,” *J. Opt. Soc. Amer.*, vol. 57, no. 9, pp. 1105–1114, 1967.
2. J. Tessendorf *et al.*, “Simulating ocean water,” *Simulating Nat. Realistic Interactive Techn. SIGGRAPH*, vol. 1, no. 2, p. 5, 2001.
3. H. Wang, “Proving theorems by pattern recognition—II,” *Bell Syst. Tech. J.*, vol. 40, no. 1, pp. 1–41, 1961.
4. E. Veach, “Robust Monte Carlo methods for light transport simulation,” Ph.D. dissertation, Stanford Univ., Stanford, CA, USA, 1998.
5. H. Jensen, “Global illumination using photon maps,” in *Eurograph. Rendering Workshop*, vol. 96, 1996, pp. 21–30.
6. H. Igehy, “Tracing ray differentials,” in *Proc. SIGGRAPH*, 1999, pp. 179–186.
7. M. Toksvig, “Mipmapping normal maps,” *J. Graph. Tools*, vol. 10, no. 3, pp. 65–71, 2005.
8. M. Olano and D. Baker, “Lean mapping,” in *Proc. Symp. Interact. 3D Graph.*, 2010, pp. 181–188.
9. J. Dupuy, E. Heitz, J.-C. Iehl, P. Poulin, F. Neyret, and V. Ostromoukhov, “Linear efficient antialiased displacement and reflectance mapping,” *ACM Trans. Graph*, vol. 32, no. 6, 2013.[Online]. Available: http://hal.inria.fr/hal-00858220
10. C. Han, B. Sun, R. Ramamoorthi, and E. Grinspun, “Frequency domain normal map filtering,” *ACM Trans. Graph.*, vol. 26, no. 3, 2007, Art. no. 28.
11. L.-Q. Yan, M. Hasan, W. Jakob, J. Lawrence, S. Marschner, and R. Ramamoorthi, “Rendering glints on high-resolution normal-mapped specular surfaces,” *ACM Trans. Graph.*, vol. 33, no. 4, 2014, Art. no. 116.
12. L.-Q. Yan, M. Hasan, S. Marschner, and R. Ramamoorthi, “Position-normal distributions for efficient rendering of specular microstructure,” *ACM Trans. Graph.*, vol. 35, no. 4, 2016, Art. no. 56.
13. J. Harvey, “Fourier treatment of near-field scalar diffraction theory,” *Amer. J. Phys.*, vol. 47, no. 11, pp. 974–980, 1979.
14. P. Beckmann and A. Spizzochino, *The Scattering of Electromagnetic Waves from Rough Surfaces*. Turkey: Pergamon, 1963.
15. T. Zirr and A. S. Kaplanyan, “Real-time rendering of procedural multiscale materials,” in *Proc. ACM SIGGRAPH Symp. Interact. 3D Graph. Games*, 2016, pp. 139–148.
16. T. Muller, M. Papas, M. Gross, W. Jarosz, and J. Novak, “Efficient rendering of heterogeneous polydisperse granular media,” *ACM Trans. Graph.*, vol. 35, no. 6, pp. 168, 2016.
17. S. Werner, Z. Velinov, W. Jakob, and M. Hullin, “Scratch iridescence: Wave-optical rendering of diffractive surface structure,” in *Proc. Trans. Graphics SIGGRAPH Asia*, vol. 36, no. 6, Nov. 2017, doi: 10.1145/3130800.3130840.
18. B. Raymond, G. Guennebaud, and P. Barla, “Multi-scale rendering of scratched materials using a structured SV-BRDF model,” *ACM Trans. Graph.*, vol. 35, no. 4, 2016, Art. no. 57.

Contact department editor Jim Foley at foley@cc.gatech.edu.