

The Shading Probe: Fast Appearance Acquisition for Mobile AR

Dan A. Calian^{*,†}

Kenny Mitchell[†]

Derek Nowrouzezahrai^{‡,†}

Jan Kautz^{*}

^{*}University College London

[†]Disney Research Zürich

[‡]University of Montreal




Figure 1: Augmented reality scene rendered in real-time on an iPad 3 using captured shading from a 3D printed shading probe (left).

Abstract

The ubiquity of mobile devices with powerful processors and integrated video cameras is re-opening the discussion on practical augmented reality (AR). Despite this technological convergence, several issues prevent reliable and immersive AR on these platforms. We address one such problem, the shading of virtual objects and determination of lighting that remains consistent with the surrounding environment. We design a novel light probe and exploit its structure to permit an efficient reformulation of the rendering equation that is suitable for fast shading on mobile devices. Unlike prior approaches, our *shading probe* directly captures the shading, and not the incident light, in a scene. As such, we avoid costly and unreliable radiometric calibration as well as side-stepping the need for complex shading algorithms. Moreover, we can tailor the shading probe's structure to better handle common lighting scenarios, such as outdoor settings. We achieve high-performance shading of virtual objects in an AR context, incorporating plausible local global-illumination effects, on mobile platforms.

Keywords: Augmented reality, lighting capture, radiance transfer

CR Categories: H.5.1 [Multimedia Information Systems]: Artificial, augmented, and virtual realities— [I.3.7]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture;

1 Introduction

Smartphones and tablets seamlessly combine several technologies to form a particularly compelling platform from an augmented reality (AR) standpoint: high-speed cameras, high-resolution touch-screen displays, large throughput wireless antennae, and powerful mobile processing units. Despite this technical convergence, several important problems must be resolved before AR can be exported from the laboratory setting to the day-to-day mobile user.

Lighting and shading consistency is among these challenges: virtual objects displayed atop the real-world viewport must retain a spatially- and temporally-consistent appearance with their real-world surroundings. This is a difficult problem as it typically requires an accurate and responsive estimate of the real-world lighting, as well as a suitably realistic shading model that leverages this estimate to shade the virtual objects. Finally, the real-world's shading must be modulated in order to account for the coupling of the virtual object onto the real-world (i.e. by casting shadows). Our work resolves the first two of these three challenges, stable lighting estimation and realistic shading, with negligible computational requirements suited for the constraints of mobile platforms.

Traditionally, this problem has been solved in three stages: incident illumination registration and capture, radiometric calibration, and shading. Normally, a reflective sphere is placed in the scene and registered by the AR application; then, a temporally-stable high-dynamic range (HDR) estimate of incident illumination must be computed from the imaged sphere; finally, some model of realistic light transport must be applied to the virtual objects, using the captured illumination, in order to simulate the effects of shadows and interreflections. The limitations of this model are two-fold: firstly, HDR capture on smartphone video cameras is prone to noise, requiring both costly radiometric calibration and temporal smoothing post-processes (to avoid jarring temporal “snapping” artifacts in

the final shading); secondly, all of the various realistic light transport shading models proposed in the past (see Section 2) remain prohibitively costly given the computational constraints of mobile applications. We adopt a hybrid solution to this problem.

We present a specially-designed 3D-printed *shading probe* to directly capture diffuse global illumination (and not incident illumination) for scene-consistent shading. At run-time we efficiently shade virtual objects, including shadowing and global illumination, using the shading response captured off our probe. This has several benefits: we can capture (diffuse) shading, which is not an HDR signal, in a single exposure (on a smartphone camera) and we avoid costly illumination calibration and light transport evaluation/simulation. Our shading probe’s shape can be tailored to suit different lighting scenarios (e.g., outdoor vs. indoor settings) and it is precisely this shape that induces a re-parameterization of the rendering equation that allows us to efficiently shade virtual objects in a manner similar to previous work on *precomputed radiance transfer* (PRT; see Section 2).

We demonstrate the effectiveness of our shading probe, computing realistic AR shading on an iPad, all while maintaining a consistent appearance with the surrounding real-world environment. We compare favorably to ground truth light transport simulation quality, all while outperforming the state-of-the-art in mobile PRT techniques.

2 Related Work

We outline recent advances in mobile AR, related works in consistent shading, and rendering approaches that motivate our work.

Before virtual objects can be placed and shaded in AR, a world-space coordinate frame must be computed. Existing solutions either use markers placed in the real scene, e.g., LEDs, colored dots, printed patterns, or use markerless vision-based calibration [Schmalstieg and Hollerer 2013]. We employ a natural feature-based registration approach by Qualcomm (Vuforia) and instead focus on simplifying lighting capture and realistic shading of virtual objects once they have been placed in the scene.

Image-based lighting has been used in AR [Nowrouzezahrai et al. 2011], with incident lighting captured off e.g. mirror light probe spheres [Debevec 2002] or unshadowed diffuse markers [Debevec et al. 2012]. HDR incident light can be captured with multiple low dynamic range (LDR) [Debevec and Malik 2008] exposures, multi-orientation photographs [Szeliski and Shum 1997], specialized “single-shot” light probes [2012], fish-eye cameras [Knecht et al. 2010], videos [Diverdi et al. 2008], or via simulation [Goesele et al. 2003]. We instead directly capture *shading* on a specially-structured diffuse probe, simplifying both lighting capture and shading: we completely avoid the need to process the captured data (e.g., performing expensive radiometric calibration, temporally smoothing, or basis projections), and the probe’s shape induces a natural basis for efficient shading of virtual objects.

Given the captured incident lighting, several approaches can compute the shading on virtual objects. Fournier et al. [1992] roughly matches external lighting with manually placed virtual light sources and precomputes global illumination solutions to transfer diffuse interreflections onto simple proxy geometries. We do not couple lighting between virtual and real objects [Jacobs et al. 2005; Knecht et al. 2010; Knecht et al. 2012], but instead use shading information from the real scene (via our shading probe) to approximate shading on arbitrary virtual objects. Nowrouzezahrai et al. [2011] capture incident illumination from a mirror sphere and compute its spherical harmonics (SH) representation for use with complex shading models at run-time. Knecht and colleagues [2011] similarly capture the incident illumination using a fish-eye camera and approxi-

mate shading using imperfect shadow maps [Ritschel et al. 2008]. None of these techniques scale to mobile platforms. In fact, most mobile AR solutions rely on simple shading models, often completely neglecting that shadows and interreflections are essential to the realistic “placement” of virtual objects in AR.

PRT projects incident light and light transport operators onto basis representations suitable for fast relighting at run-time [Sloan et al. 2002; Ramamoorthi 2009]. We are motivated by *modular radiance transfer* [Loos et al. 2011; Loos et al. 2012], where these operators are warped from simple shapes onto more complex geometries at run-time: we also capture shading information from the scene, forming a novel basis induced by the geometry of our probe, to very efficiently compute shading on our virtual objects. The shape of our shading probe (see Figure 2) can be tailored to better sample specific lighting conditions; for example, one can employ a different shading probe depending on whether the AR scenario will be indoors or outdoors. This tailoring of the shading probe can also be used to trade between accuracy, compute time and capture time.

3 The Shading Probe




Figure 2: Shading probe designs with spherical partitions created by piecewise-constant basis functions; these are, from left-to-right: 6D, 14D and 14D tailored for indoor and outdoor AR respectively.

We propose an illumination acquisition method that enables real-time global illumination rendering of complex AR scenes which are seamlessly integrated within the real environment through the use of specially designed light probes.




Figure 3: The left image is a photograph of a 3D printed 6D shading probe. The right image shows all the six acquired shading samples from the probe. Each overlaid colored disk represents such a sample, with its center showing the actual shading value and its outline indicating current from-camera visibility.

We introduce a family of shading probe designs which enable useful sampling of *shadowed shading* using an adjustable discretization. Such a probe captures actual shading in the form of shading values at n flat patches located on the probe, each of which only receives light from a particular spherical section. These sections constitute the discretization of spherical visibility into disjoint convex partitions. For an intuitive visualisation of the proposed family of probe designs please see Fig. 2. The acquired shading is then converted for use in a standard PRT framework [Sloan et al. 2002], but with piecewise constant spherical basis functions that correspond to the partitions (more in Section 3.2).




Figure 4: General comparison, from left-to-right, between: 14D, 14D-Indoor, 14D-Outdoor, 25 coefficients SH and a 110D probe. Bottom-left corner of each image shows the reconstruction of incident radiance of the corresponding method.

3.1 Shading Probe Design

We support an arbitrary number of spherical partitions in our shading probe design. In the simplest case, we uniformly distribute k directions over the sphere and use the spherical Voronoi diagram to define the spherical partitions. The most straightforward design has six faces, which correspond to the projection of cube faces onto the sphere (Fig. 2a). The actual corresponding physical probe is shown in Fig. 1 and 3a. Throughout the article, we refer to individual shading probes by their partitioning dimensionality; i.e. 14D refers to an approximately-uniformly partitioned probe into 14 spherical zones.

Indoor and Outdoor Optimization If it is known that the probe will be used indoors, where there is no single dominant light source, then we alternatively subdivide the sphere using a denser distribution on the upper hemisphere. I.e., we subdivide more finely at the top, where most lighting is likely to come from and less so toward the bottom (Fig. 2c). For outdoor use, we propose to customize the probe according to the predicted sun path (for a given day and location). In this case, we discretize more finely along the sun path and coarsely everywhere else, as the dominant illumination will be the sun (Fig. 2d).

Geometry and Manufacturing We physically build the shading probe, using 3D printing, from two main parts: the black outer shell, thus colored to suppress indirect light bounces, and the kernel printed in a, conveniently white, diffuse material whose faces capture the required constant basis-convolved shading samples. The geometry of both parts is automatically generated starting from a description of a probe’s partitioning as directions.

3.2 Acquisition and Rendering

After the 3D printed probe is registered using a 2D printed marker, the acquisition process requires the user to rotate the camera around the probe (Fig. 3). The n shading samples, one for each partition, are automatically captured from the faces of the probe’s kernel when they are deemed visible on the physical probe. Visibility is determined through raycasts to the virtual probe. The radiance samples are linearised by inverting the camera gamma before rendering, and converted back to gamma space for display.

The acquired shading is sampled only very coarsely. For our simplest shading probe, we acquire only six different shading values – one for each of the six partitions. The shading value S_i for direction i corresponds to: $S_i = L(\mathbf{n}_i, V_i) = \frac{\rho_p}{\pi} \int_{\Omega} L_{in}(\omega) V_i(\omega) (\mathbf{n}_i \cdot \omega)^+ d\omega = \frac{\rho_p}{\pi} \int_{P_i} L_{in}(\omega) (\mathbf{n}_i \cdot \omega)^+ d\omega$, where $L_{in}(\omega)$ is the incident radiance, V_i is the visibility for partition P_i and \mathbf{n}_i is the normal of the patch at partition P_i . Assuming $L_{in}(\omega)$ is constant across the i -th partition, we can denote it by L_{in_i} and move it outside the integral; re-arranging the equation yields the value L_{in_i} . We define the spherical piecewise-constant basis function $b_i(\omega)$ (normalized to be orthonormal) which corresponds to the spherical partition P_i ; L_{in_i} is then the coefficient for the basis function b_i . This enables us to use

the acquired shading in a standard PRT framework [2002]. To perform PRT using the basis induced by a specified probe we preprocess the virtual scenes to precompute global illumination transport coefficients and store them in texture maps or at vertices. At runtime, AR scenes are shaded in real-time by computing a dot product of the captured shading samples and the transport coefficients.

Envisioning the general idea of capturing implicitly computed basis projections of incident illumination for other commonly used spherical basis sets, such as SH or Haar wavelets, would not be possible using a single physical probe. For example, three-bands SH would require nine physical probes to be manufactured and registered, while sampling the shading points from the probes’ surfaces would be difficult as they would be very occluded.

4 Results

Figures 4 and 5 contain comparisons between different variants of our approach contrasted with spherical harmonics PRT. As we inherit from PRT, our shading probe approach handles diffuse inter-reflections as standard PRT methods, and this can be observed in the results.

As expected, color reproduction suffers for the shading probes from the imposed choice of a piecewise constant spherical basis set. The sun-path probe suffers most, as most of its partitions are by definition aligned with the sun path and thus it only samples the rest of the environment very sparsely. The center image in each figure clearly shows that high frequency shadows are recovered by the 14D sun-path optimized shading probe; moreover, the sun shadow fidelity obtained by this probe is matched only by the heavily-discretized 110D probe, while all the other methods (including the other 14D probes and the 25 coefficients SH baseline) only produce smoothly varying shadows. The bottom scene in Figure 5, of a dragon illuminated by an indoor probe, exhibits small variations across the renderings. However, a closer look (zoomed-in segments) reveals higher frequency shadows on the upper part of the dragon for the indoor-optimized probe when compared to the equivalently partitioned 14D probe.

5 Limitations and Future Work

While acquiring shading from the shading probe, occasional slight misalignments happen due to strong wind, imperfect extrinsic camera calibration or accentuated occlusion of the printed marker. As such, in future work we desire to experiment with other tracking solutions such as 3D model-based tracking for anisotropic probes, or by embedding markers within the unsampled probe geometry.

In the limit, a shading probe’s surface becomes specular and it could be used to render surfaces with view-dependent reflectance distribution functions (i.e. specular or glossy). However, staying within reasonable memory and fabrication constraints a limitation of the current probe design is that it only handles diffuse light transport, and we believe that overcoming this limitation poses an interesting problem for further investigation.




Figure 5: Comparison between the same methods and ordering as in Fig. 4 using different scenes: top scene shows thin poles producing high-frequency shadows; while the bottom scene shows a dragon illuminated by an interior light probe. Green-outlined insets highlight areas of interest.

6 Conclusion

Compelling and immersive AR remains a notoriously difficult problem since the human perceptual system is extremely sensitive to even the subtlest discrepancies in spatial consistency or appearance present between virtual and real objects. Capturing shading directly in the LDR domain using a novel shading probe allows us to bypass radiometric calibration, and augment reality with seamlessly integrated virtual objects. Furthermore, our shading method produces temporally coherent results with registration, capture and rendering running in real-time even on consumer-level mobile devices.

Acknowledgments This work was in part funded by Disney Research Zürich, the VEIV Centre at University College London, the Engineering and Physical Sciences Research Council and the EU FI-PPP project FI-CONTENT 2. We thank the anonymous reviewers for providing useful feedback towards improving the paper.

References

- DEBEVEC, P., AND MALIK, J. 2008. Recovering high dynamic range radiance maps from photographs. In *ACM SIGGRAPH 2008 Courses*, ACM, 31.
- DEBEVEC, P., GRAHAM, P., BUSCH, J., AND BOLAS, M. 2012. A single-shot light probe. In *SIGGRAPH Talks*.
- DEBEVEC, P. 2002. Image-based lighting. *IEEE Comput. Graph. Appl.* 22, 2 (Mar.), 26–34.
- DIVERDI, S., WITHERT, J., AND HLLERERT, T. 2008. Envisor: Online environment map construction for mixed reality. In *Proc. IEEE VR 2008*.
- FOURNIER, A., GUNAWAN, A. S., AND ROMANZIN, C. 1992. Common illumination between real and computer generated scenes. Tech. rep., University of British Columbia.
- GOESELE, M., GRANIER, X., HEIDRICH, W., AND SEIDEL, H.-P. 2003. Accurate light source acquisition and rendering. In *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 621–630.
- JACOBS, K., NAHMIA, J.-D., ANGUS, C., MARTINEZ, A. R., LOSCOS, C., AND STEED, A. 2005. Automatic generation of consistent shadows for augmented reality. In *Graphics Interface*, 113–120.
- KNECHT, M., TRAXLER, C., MATTAUSCH, O., PURGATHOFER, W., AND WIMMER, M. 2010. Differential instant radiosity for mixed reality. In *IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, 99–107.
- KNECHT, M., DÜNSER, A., TRAXLER, C., WIMMER, M., AND GRASSET, R. 2011. A framework for perceptual studies in photorealistic augmented reality. In *Proceedings of IEEE VR Workshop on Perceptual Illusions in Virtual Environments*, 27–32.
- KNECHT, M., TRAXLER, C., MATTAUSCH, O., AND WIMMER, M. 2012. Reciprocal shading for mixed reality. *Computers & Graphics* 36, 7, 846–856.
- LOOS, B. J., ANTANI, L., MITCHELL, K., NOWROUZEZHAI, D., JAROSZ, W., AND SLOAN, P.-P. 2011. Modular radiance transfer. *ACM Trans. Graph. (ACM SIGGRAPH Asia)* 30, 6.
- LOOS, B., JAMES NOWROUZEZHAI, D., JAROSZ, W., AND SLOAN, P.-P. 2012. Delta radiance transfer. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*.
- NOWROUZEZHAI, D., GEIGER, S., MITCHELL, K., SUMNER, R., JAROSZ, W., AND GROSS, M. 2011. Light factorization for mixed-frequency shadows in augmented reality. In *IEEE International Symposium on Mixed and Augmented Reality*, 173–179.
- RAMAMOORTHY, R. 2009. Precomputation-based rendering. *Found. Trends. Comput. Graph. Vis.* 3, 4 (Apr.), 281–369.
- RITSCHEL, T., GROSCH, T., KIM, M. H., SEIDEL, H.-P., DACHSBACHER, C., AND KAUTZ, J. 2008. Imperfect shadow maps for efficient computation of indirect illumination. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 27, 5, 129:1–129:8.
- SCHMALSTIEG, D., AND HOLLERER, T. 2013. *Augmented Reality: Theory and Practice*. Addison-Wesley.
- SLOAN, P.-P., KAUTZ, J., AND SNYDER, J. 2002. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 21, 3, 527–536.
- SZELISKI, R., AND SHUM, H.-Y. 1997. Creating full view panoramic image mosaics and environment maps. In *SIGGRAPH '97*, 251–258.