An Approximate Shading Model for Object Relighting

Zicheng Liao[1,2], Kevin Karsch[1], David Forsyth[1]
[1]University of Illinois at Urbana-Champaign. [2]Zhejiang University.

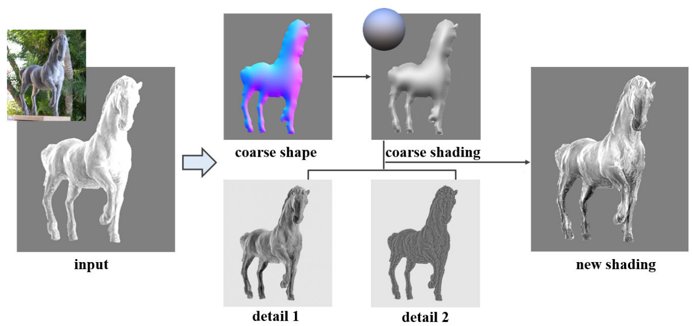

Figure 1: Given an image fragment with albedo-shading decomposition (left, albedo is omitted in this illustration), we build a coarse shape from contour; we then derive two shading detail layers (*detail 1* and *detail 2*) from the shading image and the coarse shape (middle). Our approximate shading model reshades the object under new illumination and produce a new shading on the right. Notice the change in both gross shading (new lighting from above) and the surface detail. The detail images are rescaled for visualization.

Figure 2: Given an existing 3D scene (left), our method builds approximate object models from image fragments (two teapots and an ostrich) and inserts them in the scene (right). Lighting and shadowing of the inserted objects appear consistent with the rest of the scene. Our method also captures complex reflection (front teapot), refraction (back teapot) and depth-of-field of effects (ostrich). None of these rendering effects can be achieved by image-based editing tools, while our hybrid approach avoids the pain of accurate 3D or material modeling. *3D scene modeling credit to: Jason Clarke*.

An important task of image composition is to take an existing image fragment and insert it into another scene. This approach is appling because 3D models are difficult to build, and image fragments carry real texture and material effects that achieve realism in a data-driven manner. Pure image-based composition methods (e.g., [1, 9]) rely on artist's discretion to choose source image fragments are that have compatible shading with the target scene. A relighting method would largely expand the range of images to composite with. However, it requires known shape and material of the source object. Estimating these from single image fragment is a very challenging task. State-of-the-art algorithms, such as the SIRFS method of Barron et al. [2], still produce weak shapes and do not work well for complex materials on real data. Is there a compromise between the two spaces?

We propose such an approach by exploiting an approximate shading model. The model circumvents the shape from shading problem, yet it is flexible and can adapt to various target scene lighting conditions.

The model is inspired by two lines of work. First, it is not currently possible to produce veridical reconstructions from single images; but studies show the human visual system is tolerant of renderings that are not physical [4, 8]. Our model attempts to exploit this fact to conceal *weak reconstructions*, following the spirit of work in material editing [6] and illumination estimation [5]. Second, illumination cone theory [3] suggests an accurate shading can be expressed as a linear combination of a few components. Our method decomposes shading into 3 components - a smooth component captured by a coarse shape h, and two shading detail layers: parametric residual S_p and geometric detail S_g. A new shading is expressed as the coarse shading plus a weighted combination of the two detail layers:

$$S(h,S_p,S_g) = shade(h,L), + w_p S_p + w_g S_g \quad (1)$$

where L is (new) illumination, w_p and w_g are scalar weights (Fig. 1).

The coarse shape produces a smoothing shading that captures directional and coarse-scale illumination effects that are critical for perceptual consistency. The shape is reconstruction based on contour constraints (shape from contour), which is stable and robust to generic views. The two detail layers account for visual complexity. The parametric residual S_p is residual of fragment shading estimate (recovered from image by color Retinex) and a parametric shading image, which is the shading of the coarse shape by a parametric illumination fit to the shading estimate. The geometric detail is computed by a patch-based filter introduced in Liao et al. [7]. The two detail layers encode mid and high frequencies of the shading left out by the smooth shading respectively. Intuitively, the mid frequency shading detail corresponds to object-level features (silhouettes, creases and folds, etc.), and the high frequency shading detail correspond to material property. Notice the image-based composition in the shading model is not physically-based. In practice, however, it yields remarkably good results.

With this approximate shading model, we build an object relighting system that allows an artist to select an object from an image and insert it into a target scene. The target scene can be existing 3D scene, or built from an image. The artist then selects an object model and places into the scene. Though rendering and detail composition, the system can adjust illumination on the inserted object so that it appears as it belongs in the scene (Fig. 2). Our quantitative evaluation and extensive user study suggest our method is a promising alternative to existing methods of object insertion.

This is an extended abstract. The full paper is available at the Computer Vision Foundation webpage.

[1] A. Agarwala, M. Dontcheva, M. Agrawala, S. Drucker, A. Colburn, B. Curless, D. Salesin, and M. Cohen. Interactive digital photomontage. *ACM Trans. Graph.*, 23(3):294–302, 2004.

[2] Jonathan T Barron and Jitendra Malik. Color constancy, intrinsic images, and shape estimation. *ECCV*, 2012.

[3] P. N. Belhumeur and D. J. Kriegman. What is the set of images of an object under all possible illumination conditions? *IJCV*, 1998.

[4] Patrick Cavanagh. The artist as neuroscientist. *Nature*, pages 301–307, 2005.

[5] Kevin Karsch, Varsha Hedau, David Forsyth, and Derek Hoiem. Rendering synthetic objects into legacy photographs. In *ACM Trans. Graph (SIGGRAPH Asia)*, volume 30, pages 157:1–157:12, 2011.

[6] Erum Arif Khan, Erik Reinhard, Roland W. Fleming, and Heinrich H. Bülthoff. Image-based material editing. *ACM Trans. Graph.*, 25(3): 654–663, July 2006.

[7] Zicheng Liao, Jason Rock, Yang Wang, and David Forsyth. Non-parametric filtering for geometric detail extraction and material representation. In *CVPR*, pages 963–970, Washington, DC, USA, 2013. IEEE Computer Society. ISBN 978-0-7695-4989-7.

[8] Y. Ostrovsky, P. Cavanagh, and P. Sinha. Perceiving illumination inconsistencies in scenes. *Perception*, 34:1301–1314, 2005.

[9] Patrick Pérez, Michel Gangnet, and Andrew Blake. Poisson image editing. *ACM Trans. Graph.*, 22(3):313–318, July 2003.