

Optimal Illumination for Image and Video Relighting

Francesc Moreno-Noguer

Shree K.Nayar

Peter N.Belhumeur

Department of Computer Science, Columbia University

It has been shown in the literature that image-based relighting of scenes with unknown geometry can be achieved through linear combinations of a set of pre-acquired reference images. Since the placement and brightness of the light sources can be controlled, it is natural to ask: what is the optimal way to illuminate the scene to reduce the number of reference images that are needed?

In this work we show that the best way to light the scene (i.e., the way that minimizes the number of reference images) is not using a sequence of single, compact light sources as is most commonly done, but rather to use a sequence of lighting patterns as given by an *object-dependent* lighting basis. While this lighting basis, which we call the optimal lighting basis (OLB), depends on camera and scene properties, we show that it can be determined as a simple calibration procedure before acquisition, through the SVD decomposition of the images of the object lighted by single light sources (Fig. 1).

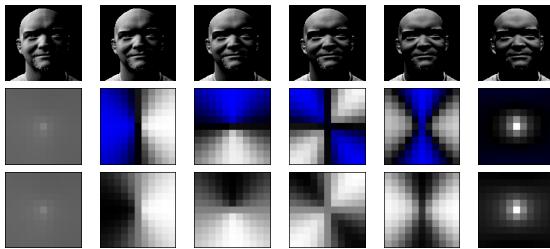


Figure 1: Computing the optimal lighting basis using SVD. First row: Images of the object illuminated by a single light source in different positions. Second row: Lighting patterns from the optimal lighting basis, containing both positive values, shown in grey, and negative values, shown in blue. Third row: Offset and scaling of the optimal lighting basis in order to make all its values positive.

We demonstrate with experiments on real and synthetic data that the optimal lighting basis significantly reduces the number of reference images that are needed to achieve a desired level of accuracy in the relit images. In particular, we show that the scene-dependent optimal lighting basis (OLB) performs much better than the Fourier lighting basis (FLB), Haar lighting basis (HaLB) and spherical harmonic lighting basis (SHLB).

In Fig. 2 we show some reconstructed images of synthetic objects which have been illuminated by SHLB and OLB. Observe how when we reconstruct from images illuminated by OLB, the error is significantly smaller. In Fig. 3 we plot the gains of the optimal lighting basis with respect the other basis, as a function of the number

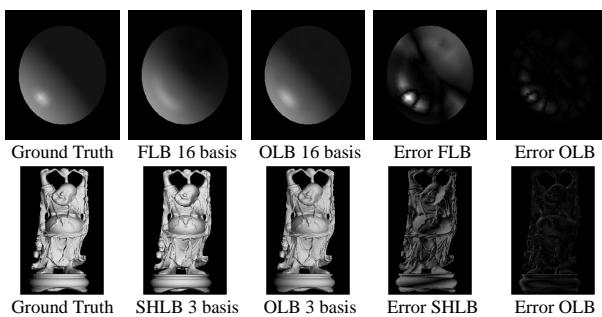


Figure 2: Examples of reconstructed images and reconstruction errors, for different lighting basis. Note that OLB performs much better.

of basis images used, and for a set of four experiments (relighting of a sphere, a face, a buddha statue, and a dragon). For any given number of optimal lighting basis images, the corresponding number of images of any other lighting basis that are needed to achieve the same reconstruction error equals the gain value. For instance, in the ‘buddha’ experiment instead of 6 optimal basis images, we will need to use $6 \times 1.8 \approx 11$ SHLB images, $6 \times 1.5 \approx 9$ FLB images or $6 \times 2.3 \approx 14$ HaLB images.

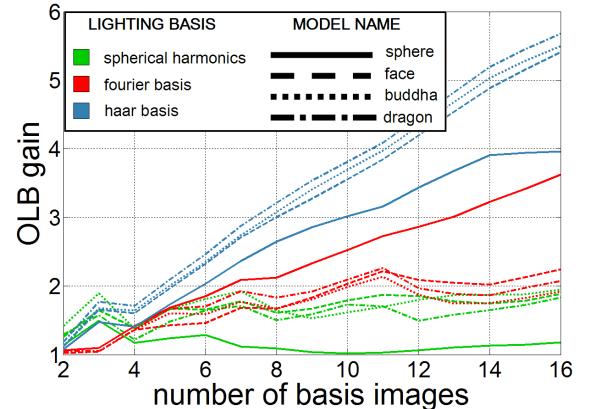


Figure 3: Gains of the OLB with respect all the other lighting basis, (for a set of 4 experiments), plotted as a function of the number of basis images used.

This reduction in the number of needed images is particularly critical in the problem of relighting in video, as corresponding points on moving objects must be aligned from frame to frame during each cycle of the lighting basis. We show, however, that the efficiencies gained by the optimal lighting basis makes relighting in video possible using only a simple optical flow alignment. Furthermore, in our experiments we verify that although the optimal lighting basis is computed for an initial orientation of the object, the reconstruction error does not increase noticeably as the object changes its pose along the video sequence.

We have performed several relighting experiments on real video sequences of moving objects, moving faces, and scenes containing both. In each case, although a single video clip was captured, we are able to relight again and again, controlling the lighting direction, extent, and color. Fig. 4 shows some frames of one of these sequences.

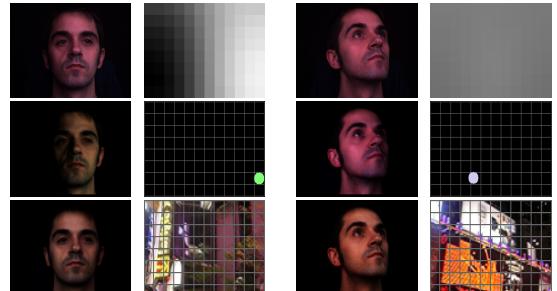


Figure 4: Two frames of a video sequence, illuminated with the optimal lighting basis (first row), and relighted with a point light source (second row) and with an environmental light (third row).