Figure S1. Here we show an example image from the experiment in Section 5.1.

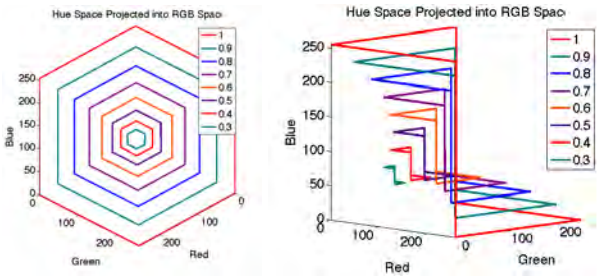

Figure S2. The space of hues with a constant saturation and value lies on a manifold in RGB space. This manifold is along a cycle around a few edges of a cube. Above we show 2 views of the same plot to show the shape of the manifolds. As the amount of light reduces, or the saturation and value of HSV measurements reduces, the measurable space in RGB gets smaller as well.

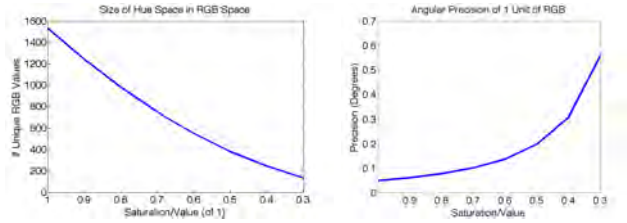

Figure S3. As the amount of light in a scene decreases, so too does the space of possible measurements of the color of a lenticular array. On the left, we show the number of unique RGB measurements for an 8-bit camera. On the right, we show how this translates to angular precision. An 8-bit camera has a maximum theoretical precision of 0.05 degrees, but would more realistically be limited to 0.1 degrees for moderately illuminated scenes.

A. Anchor Point Permutations

In Section 5.1, we explored how the anchor point localization affected HRF prediction accuracy. In that experiment, we randomly moved the 4 anchor points of a lenticular array by 8 pixels in a random direction. Figure S1 shows the anchor points used in calibrating the HRF as blue crosses and the random perturbation as green crosses.

B. The Minimum Measurable Angle Change

The angular precision resolvable of a lenticular array in an image is limited by the color sensitivity of the camera and color precision of color printing. In this section, we analyze the theoretical angular limits of a system imaging a lenticular array using an 8-bit RGB camera. The implemented lenticular arrays have their backplane textures printed at maximum saturation and value. The saturation/value of the lenticular array's appearance to the camera, however, is also determined by the amount of light in the scene and camera properties such as exposure. We explore how the amount of light affects the angular measurement precision in this section.

The set of RGB values corresponding to the hue wheel for a given saturation/value level lie on a 2d manifold in RGB space. This manifold represents the set of RGB measurements a camera would take of a lenticular array in any orientation. In Figure S2, we show two views of these manifolds for various levels of saturation and value. The manifolds create a cycle along the sides of a cube aligned with RGB space. Interestingly, as the amount of light (saturation and value) goes down, the set of RGB values corresponding to the hue wheel gets smaller.

On the left of Figure S3, we show the number of unique 8-bit RGB triplets for each manifold, or each saturation/value level. The set of 1-to-1 view points of the lenticular array (76 degree range of incident angles) is mapped to the hue wheel and therefore to the RGB manifold. Thus, our angular precision is 76 / the number of unique rgb values. On the right of Figure S3, we show this angular precision for each level of saturation/value. In the best possible case with maximum saturation and value, an 8-bit RGB camera is able to resolve the angle of a lenticular array at a precision of 0.05 degrees. However, at 0.3 saturation/value, the precision drops to 0.55 degrees. For the experiments in Section 5.1, we note that the hue measurements have a mean saturation and value of ≈ 0.7 and ranged from ≈ 0.5 to ≈ 0.9. Therefore, we do not believe that the angular error of 0.25 degrees induced by moving anchor points is due to the inherent precision limitations of the 8-bit RGB camera imagine the lenticular arrays.

The angular precision achievable by a camera can be greatly improved by moving to a larger color representation. In Figure S4, we show the same experiment as before, but for a 12-bit camera. With 16 times more possible values for a single color channel versus an 8-bit camera, the number of unique RGB values for the color wheel and the angular precision both improve by an order of magnitude. Therefore, in future work, we plan to explore camera calibration and pose estimation using higher bit precision cameras.

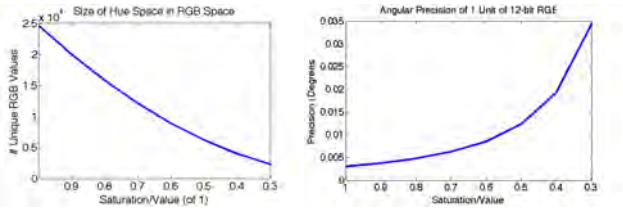

Figure S4. We repeated the experiment depicted in Figure S3, but for a camera that can capture 12-bit RGB images. As the camera can now measure color with 16 times more values, we achieve an order of magnitude more of uniquely measurable hues in RGB space (left) and therefore an order of magnitude smaller angular precision.

C. Additional Augmented Reality Video

We demonstrate our approach with a second AR video where the camera is static with a varying zoom, and the calibration object is being rotated randomly. In this video the zoom is achieved via a zoom lens, in contrast to the digital zoom performed in Section 5.3. In the supplementary material, we include 2 additional videos. In the first video, we overlay the wire-mesh of a box to compare dynamic focal length estimation versus static focal length estimation (just like in Section 5.3). In the second video, instead of a box wire-mesh, we overlay a 3D model of a parrot over the frames of the image. In Figure S5, we show frames of this video. Just like in the previous results, our dynamic focal length estimation ensures that the 3D model is rendered with the correct perspective, no matter the zoom level.

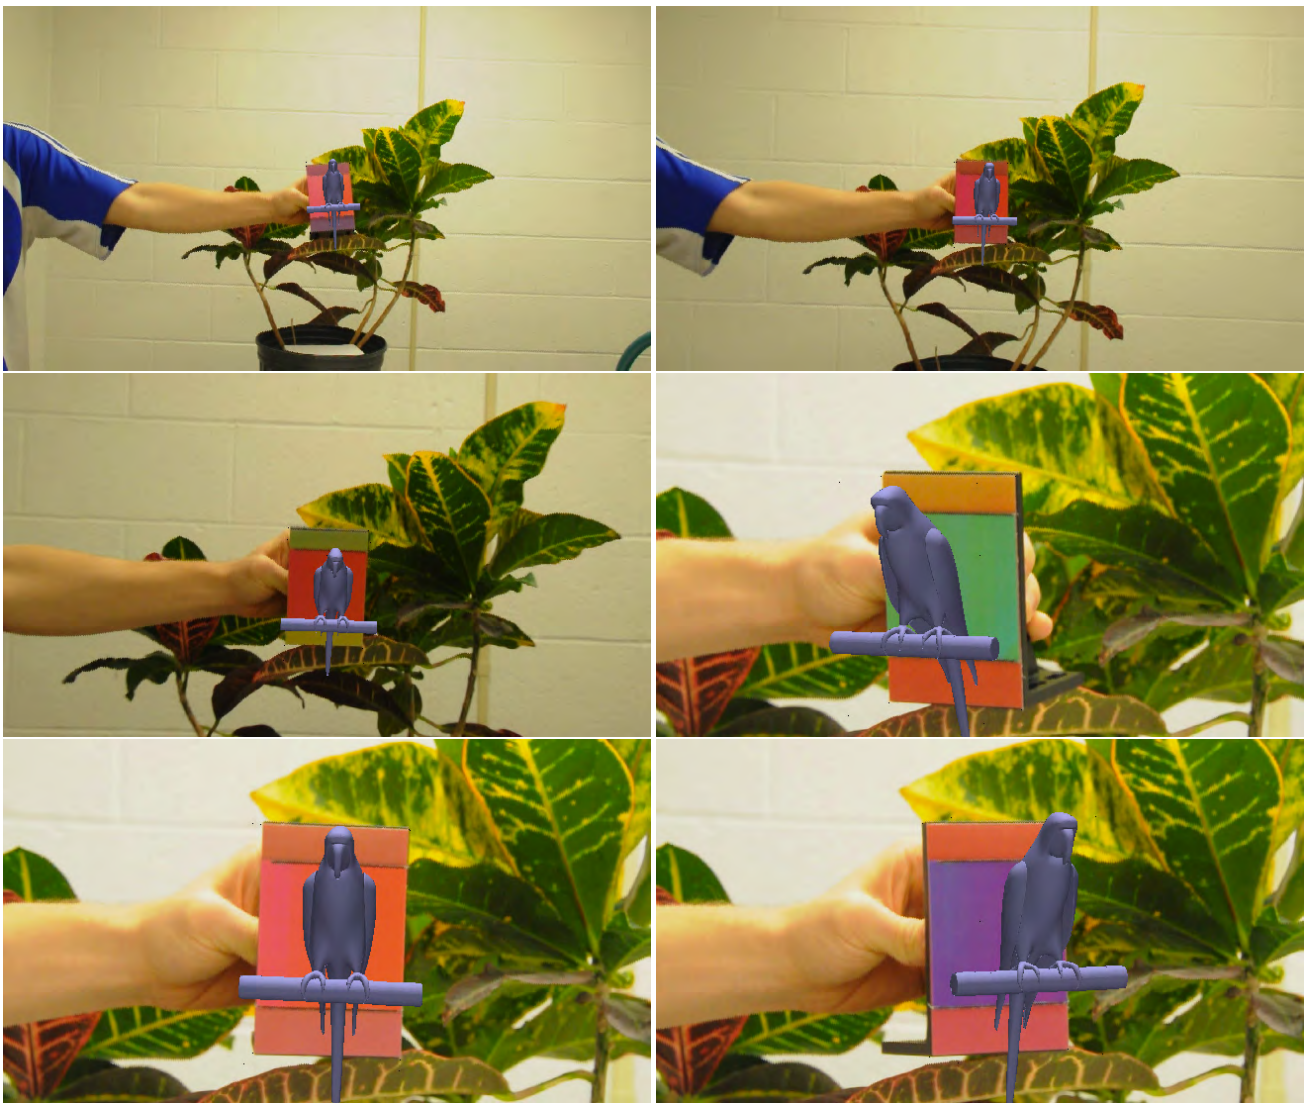

Figure S5. We use our focal length, rotation, and translation estimates to overlay a 3D model of a parrot onto the image. Because we estimation the focal length at each image, the parrot is rendered with the appropriate perspective despite very different zooms throught the video.