Priors for Large Photo Collections and What they Reveal about Cameras

Adam Hewgill

Sujit Kuthirummal, Aseem Agarwala, Dan B. Goldman, and Shree K. Nayar. Priors for Large Photo Collections and What they Reveal about Cameras. In ECCV '08: Proceedings of the 10th European Conference on Computer Vision, pages 74-87, Marseille, France, 2008.

Kuthirummal et al.'s Priors for Large Photo Collections and What they Reveal about Cameras takes advantage of the availability of large photo collections on the Internet—specifically Flickr—which contain metadata for various properties describing the camera and its settings. From this information groupings of photos are made with similar properties for statistical analysis. Using the results of the analysis they are able to discover the non-linear response function of a camera model, spatially-varying vignetting of a camera models lens settings, and discover bad pixels of the detectors of specific cameras. These calibration settings can them be used to setup a camera without using a tedious physical calibration process.

In previous work other researchers have discovered methods to find the response functions but most require large amounts of user time and access to the camera models. Another approach uses a single photo and makes assumptions about the shape of the curve which in practice often causes significant error. Techniques exist for determining vignetting but are labour intensive and very difficult to setup correctly. Specialized hardware exists but is prohibitively expensive for wide use. Bad pixels on the detector are a common occurrence and arise over time. A previous method uses Bayesian filters to discover the pixels using a set of images but the results are not definite or easy to produce.

The researchers selected a set of 5 popular cameras to use for this research: Canon S1IS, Sony W1, Canon G5, Casio Z120, and Minolta Z2. Approximately 40,000 images were collected for each type of camera spread across the various lens settings. Each camera was manually calibrated using a software tool called HDRShop to produce the ground truth curves for later experimental validation. The images from the Canon S1IS camera model were then undistorted to serve as the camera-independent priors. The EXIF metadata then allows the researchers to group the photos for each camera model based on the lens settings. With this data they compute the spatial distribution of average image luminance and the joint histogram of irradiances at neighbouring pixels and compare it with the values computed from the camera-independent priors. The differences provide the radiometric properties of each camera model. By analyzing the average irradiances of neighbouring pixels the researchers discover outliers (differing by 7 grey levels) which indicate point detector failure—i.e. bad pixels.

The key benefit of this technique is the ability to discover the calibration settings required for a specific camera model with a specific lens setup using a diverse collections of readily available Internet images. These values can be determined without requiring a complex controlled capture environment and even without physical access to the camera. This simplifies the camera configuration process by reducing the amount of work necessary. Finding ways to take advantage of large data sets to reduce the efforts required to perform certain tasks is a common research effort and can be very effective.

Results provided in Figure 3 for the response functions and in Figure 5 for the vignetting show clearly that the approach does a very good job of estimating the ground truth. In Figure 6 we see an overview of bad pixel averages for 3 camera models which is an interesting finding available from the photo sets. While not simple in execution this approach shows it is capable of producing real usable results from existing data sources.

A major limitation of the approach only allows the technique to be applied to non-SLR based cameras which don't have interchangeable lenses. Lens setting metadata must be available to group photos. Photos taken in a certain way are not able to be used or they throw off the mechanics of the technique. Specifically images, captured with a flash, edit with a software tool, or where the subject is too dominant in the center of the scene must be removed. This is a major issue as you can no longer use a random sampling of images but rather must craft the set through removal of outliers, this reduces the probability you will be able to find enough samples for a specific camera model and lens setting (which is indicated to be around 3000 photos.

Extensions of this work could go in several directions from discovering other consistent camera model properties to determining a particular photographers biases or techniques. The researchers mention that other properties such as radial distortion, chromatic aberration, and spatially varying lens softness may all be able to be determined with a similar technique. It should also be possible to extend this work to SLR type cameras. We should be able to determine lens setting based on statistical analysis which could then augment the available metadata. Specialized techniques such as the use of a fisheye lens may also be detectable across all camera models. I large set of images could be processed to find all photos which use the technique and augment their metadata accordingly.