Technical Report

UCAM-CL-TR-712
ISSN 1476-2986

Number 712

UNIVERSITY OF CAMBRIDGE

Computer Laboratory

Flash-exposure high dynamic range imaging: virtual photography and depth-compensating flash

Christian Richardt

March 2008

15 JJ Thomson Avenue
Cambridge CB3 0FD
United Kingdom
phone +44 1223 763500

http://www.cl.cam.ac.uk/

© 2008 Christian Richardt

Technical reports published by the University of Cambridge
Computer Laboratory are freely available via the Internet:

http://www.cl.cam.ac.uk/techreports/

ISSN 1476-2986

Flash-Exposure High Dynamic Range Imaging: Virtual Photography and Depth-Compensating Flash

Christian Richardt

Abstract

I present a revised approach to flash-exposure high dynamic range (HDR) imaging and demonstrate two applications of this image representation. The first application enables the creation of realistic 'virtual photographs' for arbitrary flash-exposure settings, based on a single flash-exposure HDR image. The second application is a novel tone mapping operator for flash-exposure HDR images based on the idea of an 'intelligent flash'. It compensates for the depth-related brightness fall-off occurring in flash photographs by taking the ambient illumination into account.

1 Introduction

The widespread availability of consumer digital cameras has transformed the consumer photography paradigm. From taking one or two carefully posed analogue photographs, we changed to taking hundreds or even thousands of digital photographs on a single occasion.

While this is widely regarded as a significant improvement, in at least one way this transition changed things for the worse: dynamic range. Dynamic range, in this context, is the ratio between the brightest and the darkest regions in an image.

Digital photographs only have about one tenth of the dynamic range of film-based photographs, which are approximately on a par with the simultaneous dynamic range of the human eye. Particularly in outdoor scenes, this dynamic range is easily exceeded, which makes dynamic range a precious resource in digital photography.

Modern digital cameras come with sophisticated exposure control algorithms which estimate camera settings for fitting a particular scene into the available dynamic range. Despite this, it is not uncommon for the limited dynamic range of the imaging sensor to produce under- or over-exposed photographs. This is particularly likely when using the flash.

Choosing the 'right' exposure and flash settings is a hard problem which usually involves taking a lot of very similar photographs with only slightly varying settings. This has to be done at the time of shooting, and the added flexibility is limited to choosing one of the taken photographs.

I introduce a novel technique (section 3) that effectively allows selected settings – the shutter speed and the flash intensity – to be varied freely *after* shooting the initial photographs. Based on a flash-exposure high dynamic range (HDR) image, which is constructed from the initial photographs, my technique creates 'virtual photographs' for arbitrary flash-exposure settings.

Another common issue in consumer photography is that the flash often leads to a very non-uniform scene illumination. Nearby objects are over-exposed and far objects are under-exposed.

I propose a novel approach (section 4) for compensating for the brightness fall-off with depth, to create a more uniformly illuminated image. This 'depth-compensating flash' operates on a flash-exposure HDR image and produces enhanced flash images.

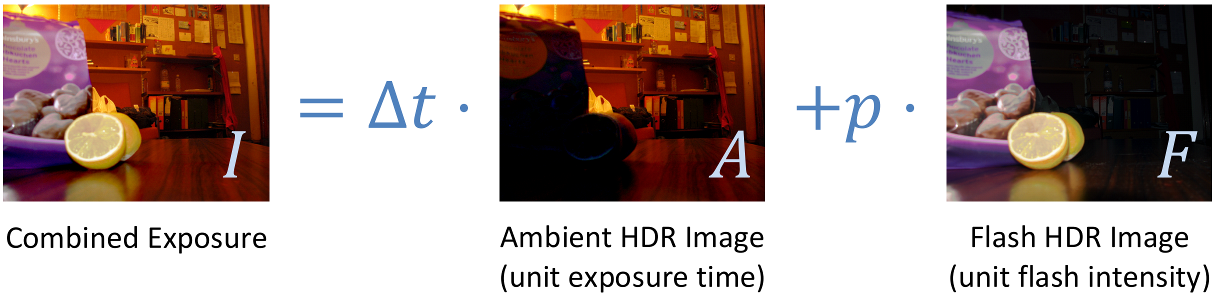

Figure 1: Visualisation of (1) showing how ambient and flash HDR images are combined.

2 Flash-Exposure HDR Imaging

Early work [4, 2] considered how to combine a number of photographs, taken in an exposure sequence, into a single HDR image. Based on this traditional approach, Agrawal et al. [1] first proposed flash-exposure HDR imaging.

Flash-exposure HDR imaging is a step towards computational photography [6] – a concept in stark contrast to conventional photography. It captures more aspects of a scene in such a way as to allow the photographer to choose selected camera settings after taking the 'photograph'. This would enable the photographer to experiment freely with these settings in an offline post-processing step after the scene has been captured.

HDR images can be considered to be one-dimensional in terms of the radiance associated with each pixel in the image. To be more precise, the value stored for each pixel is the ambient (no-flash) radiance in unit exposure time. Flash-exposure HDR imaging introduces a second value for flash intensity: the radiance due to unit flash intensity.

The flash-exposure HDR techniques presented in this paper can be used as the basis for a variety of applications, such as virtual photography and a depth-compensating flash, which are presented in sections 3 and 4 respectively.

2.1 Separability

The key observation that enables flash-exposure HDR imaging is that ambient and flash illumination in a scene are independent and can hence be separated.

Let F be the flash HDR image, or the image strictly due to a unit flash intensity. The flash-only exposure F' is then F scaled by the flash intensity p. Similarly, the ambient exposure A' is the ambient HDR image A – the image captured for unit exposure time – scaled by the exposure time Δt. The total exposure I is now given by

$$I = A' + F' = \Delta t \cdot A + p \cdot F, \qquad (1)$$

because the total amount of light incident at the image sensor is the sum of the ambient and the flash-only contributions in the scene (see figure 1).

2.2 Solution Process

In traditional HDR imaging [2], the camera response function g maps pixel values $Z(x, y)$ to logarithmic exposure $g(Z(x, y))$. In the flash-exposure case, the exposure is given by I in (1) which holds for any exposure time Δt_i and flash intensity p_j:

$$\exp\left(g\left(Z_{ij}(x, y)\right)\right) = \Delta t_i \cdot A(x, y) + p_j \cdot F(x, y)\,. \qquad (2)$$

To obtain a least squares error estimate for the ambient and flash HDR images A and F from (2), a minimum of two photographs is required. However, to increase numerical stability of the solution, a total of $n_A \times n_F$ photographs are taken at all possible combinations of n_A exposure times and n_F flash intensities (as in figure 2).

Solving the equation system derived from (2) directly would result in unacceptable colour inaccuracies and saturation artefacts in

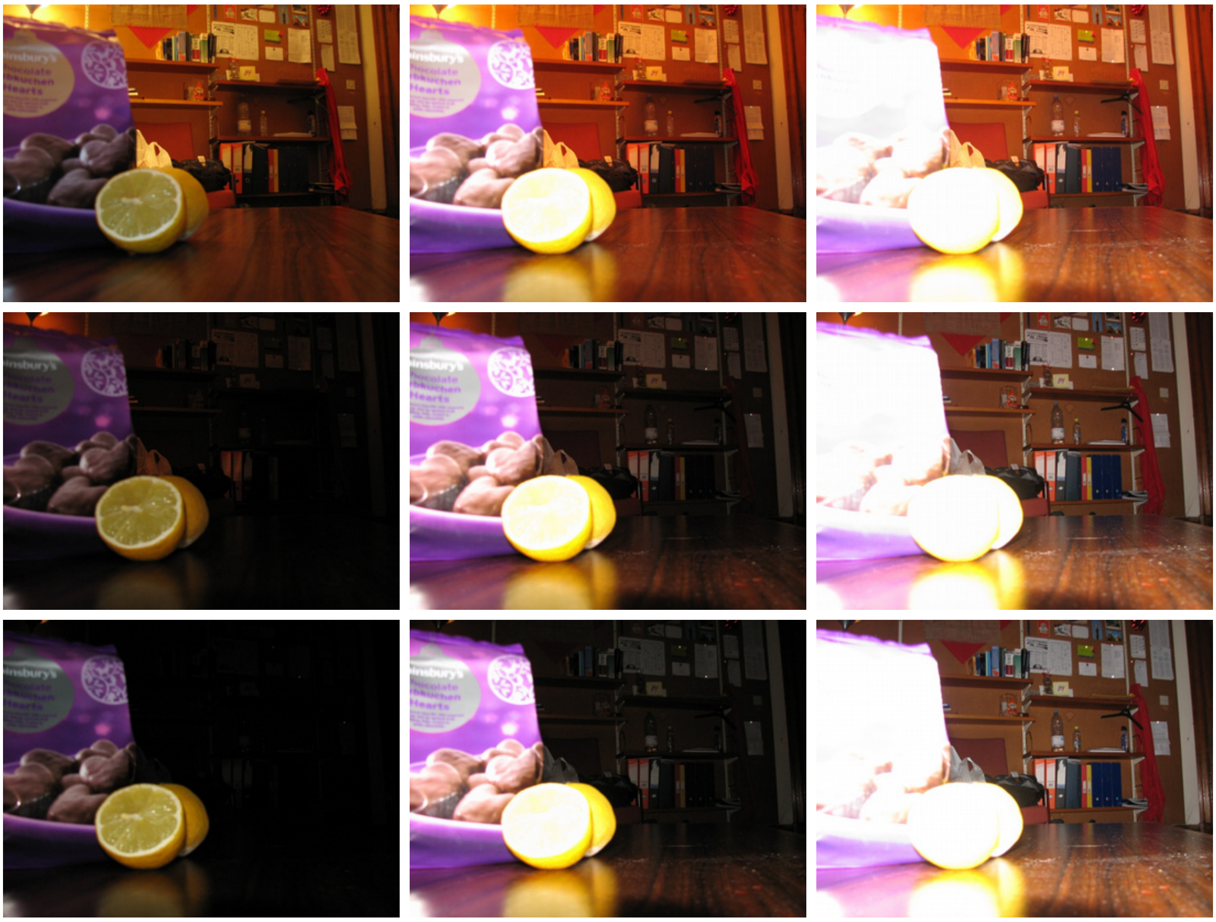

Figure 2: Input photographs for recovering a flash-exposure HDR image from three exposure times (vertical) and three flash intensities (horizontal).

regions where the input photographs are saturated. To avoid this, Agrawal et al. [1] suggested weighting equations with the signal-to-noise function of the underlying camera.

However, for general applications, a better weighted least squared error estimate can be obtained by weighting each equation by $w\left(Z_{ij}\left(x,y\right)\right)$, where $w\left(x\right)=\sin\left(\pi x\right)$ is a windowing function that reduces the effects of under- or over-exposed pixels. This considerably improves colour reproduction and reduces the aforementioned artefacts.

3 Virtual photography

Virtual photography mimics the imaging process described by a camera response curve g. This is achieved by inverting g and using it to map exposure values $I_{\Delta t,p}$ back to pixel values $Z_{\Delta t,p}$. This process can be thought of as exposing the underlying HDR image virtually.

However, contrary to conventional photography, both exposure time Δt and the flash intensity p can be varied freely and may also lie between the flash-exposure settings physically supported by the camera used to create the HDR image. Additionally, the response characteristics of any imaging process – and therefore any digital camera – can be modelled by using an appropriate response curve.

In more detail, the ambient image A and the flash image F are recombined to the exposure $I_{\Delta t,p}=\Delta t\cdot A+p\cdot F$ using (1), for the exposure time Δt and the flash intensity p. This exposure is converted to pixel values $Z_{\Delta t,p}$ by rearranging (2) as

$$Z_{\Delta t,p}\left(x,y\right)=\\g^{-1}\left(\ln\left(\Delta t\cdot A\left(x,y\right)+p\cdot F\left(x,y\right)\right)\right). \quad (3)$$

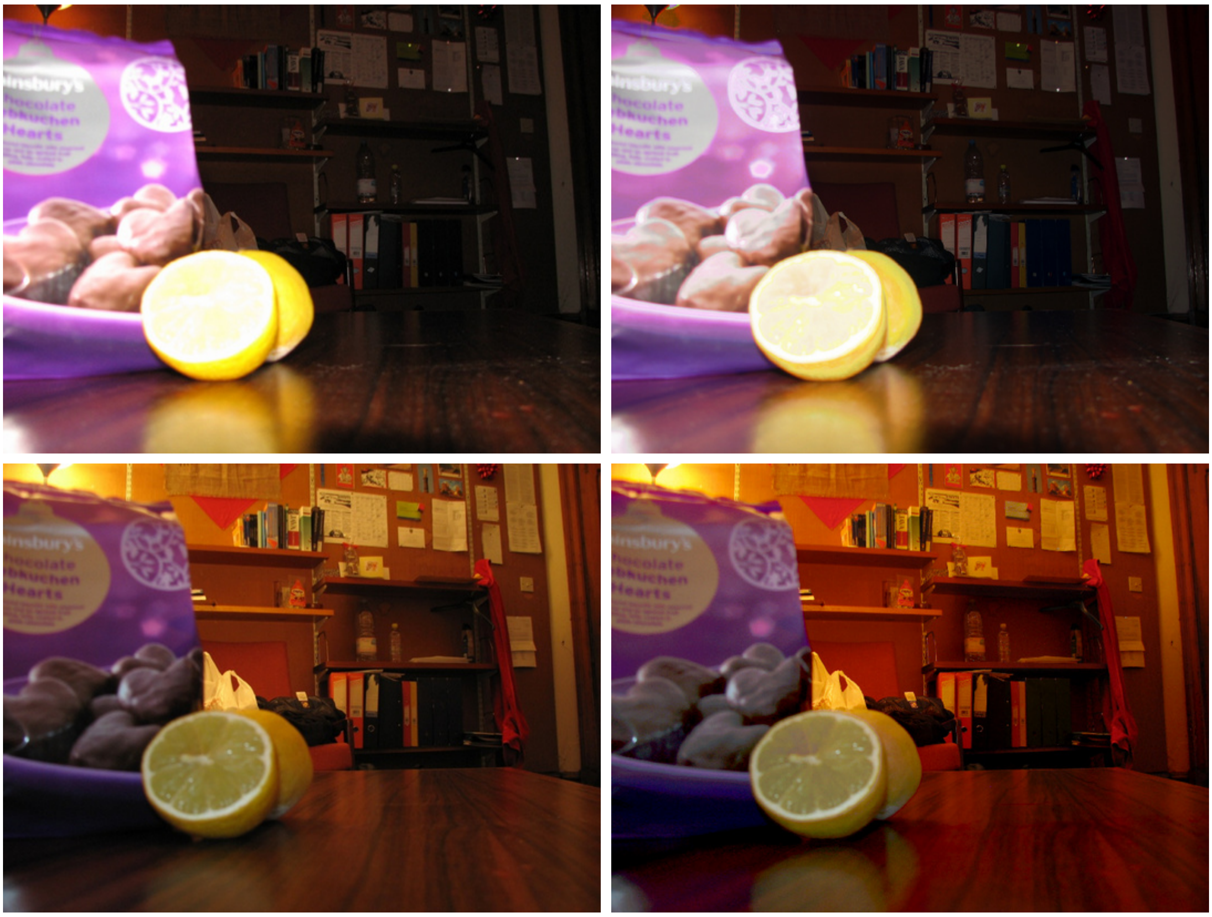

(a) Real photographs   (b) Virtual photographs

Figure 3: Comparison of (a) real and (b) virtual photographs.

Examples of virtual photographs are shown in figure 3, in which the original input photographs (a) are compared with the virtual photographs (b) obtained for the same camera parameters.

Virtual photographs can in principle be produced for arbitrary exposure time and flash intensity settings. However, if the settings lie far outside the range of parameters used in taking the original input photographs, the virtual photograph loses fidelity. In the worst case, this leads to saturation and other artefacts such as extreme amplification of residual camera noise. In order to reduce these kind of artefacts, the parameters should ideally lie within the range used in the input photograph sequence.

The combined approach of flash-exposure HDR recovery and virtual photography decouples the processes of taking the photographs and experimenting with the camera settings in a post-processing step.

4 Depth-Compensating Flash

The depth-compensating flash is a novel tone mapping operator for flash-exposure HDR images which addresses the issue that flash-induced brightness suffers a fall-off. Objects nearer to the camera tend to be over-illuminated and distant objects under-illuminated.

The algorithm presented in this section counteracts the depth fall-off using the depth-orientation map β [1] which results in a more uniform scene illumination.

The top-level design of the depth-compensating flash is outlined in figure 5. It comprises the following five stages:

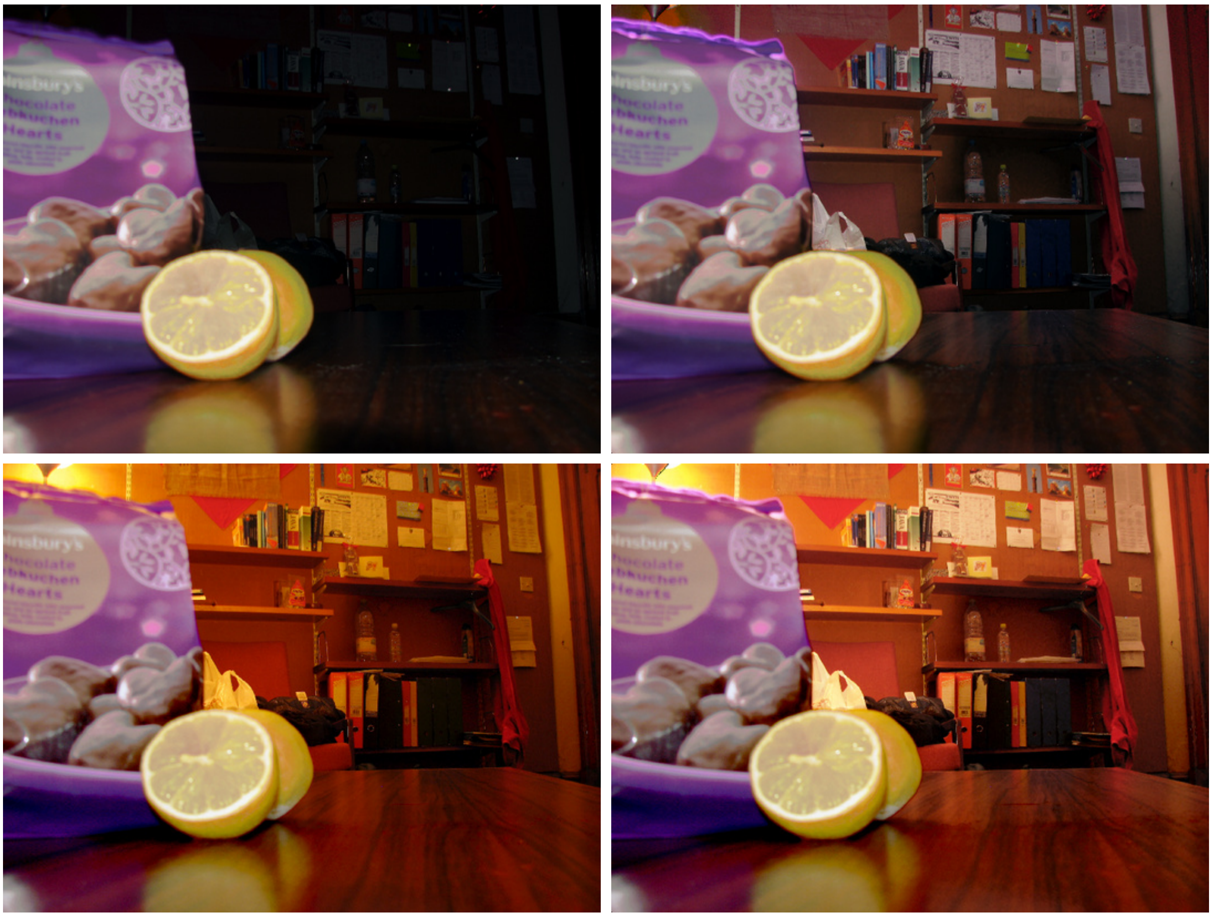

Figure 4: Comparison of virtual photographs with (right) and without (left) depth-compensating flash, for flash-only image (top) and combined image (bottom) respectively.

1. **Extract Luminance and Chromaticity**
   Since the flash predominantly affects brightness, the input images A and F are converted to the Yxy colour space to only operate on the luminance components A_Y and F_Y while leaving the chromaticity components A_{xy} and F_{xy} unchanged.

2. **Calculate Depth-Orientation Map**
   The depth-orientation map $\beta = F_Y / A_Y$ is calculated.

3. **Truncate & Filter**
   The depth-orientation map β is truncated to the range $[\beta_{min}, \beta_{max}]$ to limit the amplification and attenuation applied to the flash image in the next step. Additionally, β is filtered using a two-dimensional median filter to reduce the amount of noise contained, resulting in β'.

4. **Scale Flash Luminance**
   The flash luminance is divided by β', to compensate for the depth fall-off associated with the flash. This results in F'_Y.

5. **Recombine to Colour Image**
   Finally, the flash chromaticity F_{xy} is combined with the new flash luminance F'_Y and converted to the HDR image F' – the depth-compensated flash image.

The results of applying the depth-compensating flash are shown in figure 4. Note the improved illumination visible in the background areas on the right-hand side.

I expect that work on flash/no-flash photography [3, 5] could also be extended to flash-exposure HDR images. However, these approaches are not considered here.

5 Summary and Future Work

I presented a revised approach to flash-exposure HDR imaging and two applications based on it. Virtual photography is a technique that allows the user to generate photographs for freely chosen exposure times and flash intensities. The depth-compensating flash compensates for the brightness fall-off with depth, to create a more uniformly illuminated image.

Potential future work includes an embedded implementation of these applications on a digital camera which would calculate the flash-exposure HDR image on-board. This would lead to a lot of post-processing potential, in a similar way to 'raw' photographs.

Taking the full set of input photographs is clearly wasteful, because it comprises all possible combinations of exposure times and flash intensities. An adaptive flash-exposure sampling strategy, e.g. based on the one presented by Agrawal et al. [1], would reduce the number of required photographs considerably.

Acknowledgements

I would like to thank my proof-readers as well as Rahul Vohra, Jon Fielder and Neil Dodgson for their support and suggestions.

References

[1] A. Agrawal, R. Raskar, S. K. Nayar, and Y. Li. Removing photography artifacts using gradient projection and flash-exposure sampling. *ACM TOG (Proc. SIGGRAPH)*, 24(3):828–835, 2005.

[2] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proc. SIGGRAPH*, pages 369–378, 1997.

[3] E. Eisemann and F. Durand. Flash photography enhancement via intrinsic relighting. *ACM TOG (Proc. SIGGRAPH)*, 23(3):673–678, 2004.

[4] S. Mann and R. W. Picard. On being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures. In *Proc. IS&T Annual Conference*, pages 422–428, 1995.

[5] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama. Digital photography with flash and no-flash image pairs. *ACM TOG (Proc. SIGGRAPH)*, 23(3):664–672, 2004.

[6] R. Raskar, J. Tumblin, A. Mohan, A. Agrawal, and Y. Li. Computational photography. In *Eurographics Tutorials*, pages 557–624, 2007.

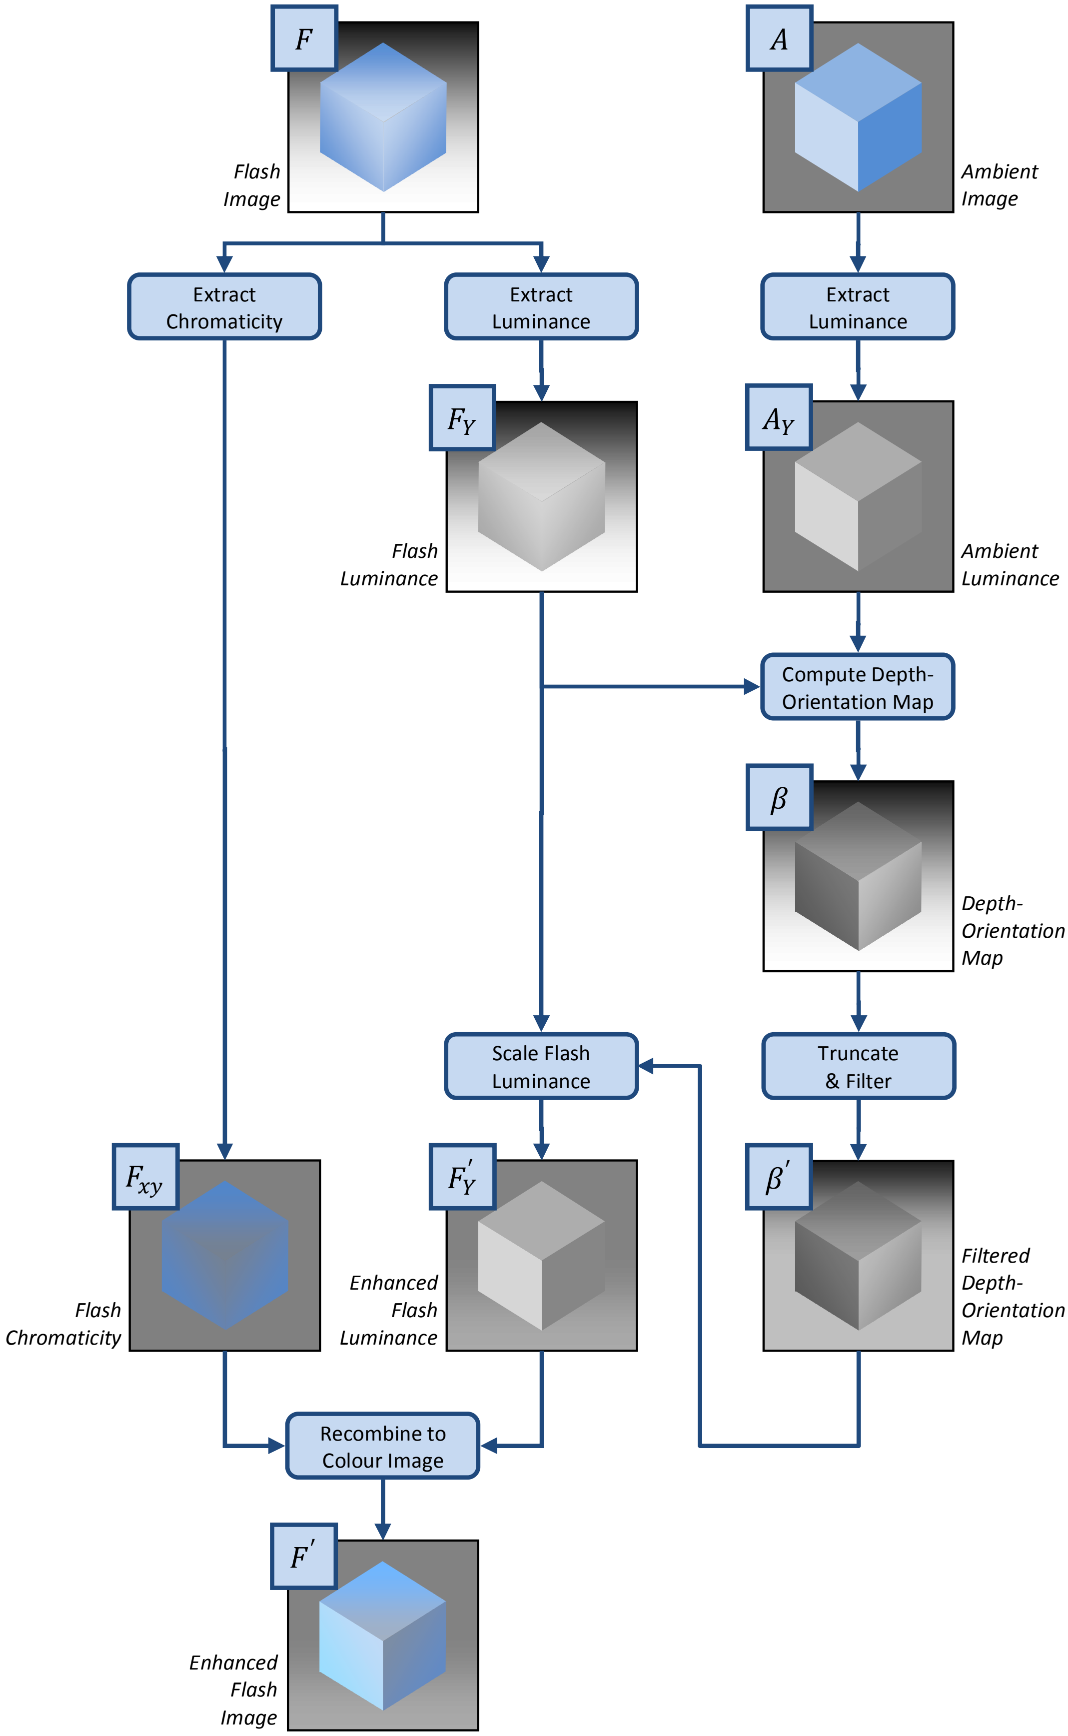

Figure 5: Top-level design of the depth-compensating flash.