

Practical Non-linear Photometric Projector Compensation

Anselm Grundhöfer
Disney Research Zürich
Clausiusstrasse 49, 8092 Zürich, Switzerland
anselm@disneyresearch.com

Abstract

We propose a novel approach to generate a high quality photometric compensated projection which, to our knowledge, is the first one, which does not require a radiometrical pre-calibration of cameras or projectors. This improves the compensation quality using devices which cannot be easily linearized, such as single chip DLP projectors with complex color processing. In addition, the simple workflow significantly simplifies the compensation image generation. Our approach consists of a sparse sampling of the projector's color gamut and a scattered data interpolation to generate the per-pixel mapping from projector to camera colors in real-time. To avoid out-of-gamut artifacts, the input image is automatically scaled locally in an optional off-line optimization step maximizing the achievable luminance and contrast while still preserving smooth input gradients without significant clipping errors.

1. Introduction

Photometric projector compensation is used in various application fields such as entertainment, cultural heritage and augmented reality. Its preparation, however, still is a laborious process. While several algorithms have been presented within the last decade which generate a compensated projection with impressive quality, they all require the devices to be, at least partially, radiometrically calibrated. This can be a cumbersome and time consuming process which reduces the overall flexibility of the system, while it also limits the compensation quality on devices such as, for example, certain DLP projectors which offer complex, non-monotonic color processing algorithms and additional primaries as well as transparent components in their color wheels [14, 13].

1.1. Motivation

The main motivation to develop a novel projector compensation algorithm has been the lack of an existing straightforward method which does not require any radio-




Figure 1. Our method computes a compensation image without any knowledge of the device's response curves. The input image (a) is projected (f) onto a textured surface (e) leading to color artifacts (b). Our method computes a compensation image (g) which reduces the errors (c). A global optimization step further minimizes saturation artifacts (d). Although the used DLP projector contains complex color processing, our compensation method is able to generate a high quality compensation image (h).

metric pre-calibration and thus can be widely applied in, for example, mobile and cheap off-the-shelf projection setups. Our method is based on the application of scattered data interpolation to describe the mapping of projector to camera pixels in a non-linear manner. We used thin plate splines (TPS) [7] to compute the mapping which, besides their relatively low computational complexity using radial basis functions (RBF), guarantee an optimally smooth transition between the captured color samples.

1.2. Background and Related Work

In this section an introduction to photometric compensation will be given and the existing methods will be summarized. The interested reader is referred to Bimber's *et al.* state of the art report [4] for details.

The main purpose of photometric compensated projections is the neutralization of non-perfectly white or textured surface reflectance. This is accomplished by using a projector-

camera system to evaluate the reflectance properties for each pixel and to calculate colors which, if projected onto the surface pigments, modulate to the expected intensity when captured by the camera. Therefore a geometric calibration is required to map projector to camera pixels which is usually carried out by projecting a series of well defined structured light patterns. Depending on the method, several patterns are projected in addition to estimate the color transformation needed to compute the compensation images. To our knowledge, all existing methods require the projector or camera to be at least partially radiometrically calibrated to enable a linear modeling of the light modulation.

The oldest algorithms focused on modeling the compensation without taking the input image content into account and mostly varied in their complexity as well as their prerequisites. In Bimber *et al.* [3], a multi-projector approach was presented in which the devices are assumed to be fully linearized and their *rgb* color channels to be completely independent. While this approach is able to also compensate for environmental illumination, most real-world setups require a modeling of the overlapping color channels to achieve a high quality compensation. In Nayar *et al.* [14], a 3×3 matrix was used to model this color channel mixing which generates a more accurate compensation. This method was integrated into a real-time feedback system in [8] to support dynamic projection surfaces. Yoshida *et al.* [22] extended this approach to a 3×4 matrix which also considers the uncontrollable, ambient illumination. In [5] it was shown that the color mixing matrix can also be separated from the spatially varying surface reflection and thus it only has to be stored once for a setup. An extended approach enabling the compensation of global, complex illumination effects such as caustics, refractions and scattering has been presented in [21]. It's based on the idea of measuring the full light transport between the projector and camera and inverting it to compute the compensation image. This approach requires, besides the device linearization, up to several hours of scanning time depending on the scene complexity which makes it quite impractical for real-world applications.

While these algorithms are able to achieve satisfying results under well calibrated conditions, image contrast will be lost and color artifacts might occur due to intensity saturation of the projector on dark surface pigments.

More recent compensation methods are focused on content dependent adaptation to increase the visual quality of the projected images by maximizing perceived contrast and luminance while still suppressing saturation artifacts at the same time. The methods optimize the projection images with different computational complexities, depending on the application scenarios. They all, however, also require a radiometrically calibrated projector-camera setup.

The first adaptive photometric compensation method was presented by Wang *et al.* [20], in which the input image

is scaled automatically by a global scaling factor until the saturation artifacts approach the per-pixel visibility threshold. The presented algorithm is only able to compensate gray scale images which constrains its applicability. This idea was extended for color images by Ashdown *et al.* In [2] they describe a compensation framework which operates in the CIE $L^*u^*v^*$ color space and applies a luminance and chrominance rescaling of the input image based on the human visual perception. This rescaling is optimized such that the visual impression after applying the photometric compensation still is close to the desired input. This algorithm was further improved in [15] enabling smoother chrominance adaptations. Another algorithm presented by Grundhöfer *et al.* [9] focused on a GPU accelerated real-time adaptation to enable the system to also work with real-time content. A sophisticated compensation method was recently presented by Aliaga *et al.* [1] which applies a globally optimized compensation using the measured light transport matrix between a camera and multiple projectors. It, however, also requires radiometrically calibrated projectors and cameras.

While those methods show that high quality compensated projections can be achieved with radiometrically calibrated devices, the latter makes the usage quite cumbersome for several reasons: Any change in the hardware settings e.g. adapting the projectors' brightness to a new setup or replacing it with a different model requires a radiometric recalibration; furthermore, some DLP projectors apply complex, multi-primary color processing and thus are difficult to accurately linearize. While projection-based illumination on public spaces becomes more and more popular, the surface colors are mostly not compensated at all because of the missing expertise and hardware to adequately calibrate the response curves of the projection system.

1.3. Our contribution

Our approach computes a compensation image by generating a non-linear per-pixel color mapping between the radiometrically uncalibrated projector-camera pair. Scattered data interpolation based on TPS is used to calculate an accurate color transformation. The advantage of the applied RBF compared to the ones used in other scattered interpolation techniques, such as Shepard's interpolation method [19] or multiquadrics [11], is its ability to smoothly interpolate between given sample points as well as to adequately extrapolate the colors. It is able to compensate for non-monotonic responses as well as inter-channel color modulations.

We present two compensation image generation methods: The first one processes images in real-time, while the second optimizes the input image off-line in a global optimization step to further minimize local clipping errors while preserving a high overall luminance and contrast.

2. Uncalibrated Photometric Projector Compensation

Instead of having to pre-calibrate the cameras and projectors, the proposed method can be applied immediately after setting up the devices. Uncalibrated projectors and cameras, however, require a per-pixel non-linear mapping function to accurately describe the unknown color transformation from the camera to the projector via the surface material. While, theoretically, this can be generated by a dense sampling to populate a 3D look-up table, this would require several millions of images to be captured and stored which is impractical considering acquisition time and memory requirements. The number of required samples can be minimized by sparsely sampling the colors and applying interpolation methods to compute the remaining values. While tri-linear interpolation obviously won't generate satisfying results, more complex interpolation methods as, for example, presented in [17], are able to minimize the errors significantly. This method, however, still requires more than 700 images to be acquired and thus is still not practical, especially if the data has to be calculated and stored for each individual pixel. To further reduce the number of color samples while nevertheless being able to achieve high quality compensation results, we apply TPS interpolation [7], which guarantees a smooth transition while interpolating between the known projected and captured sample points.

2.1. Prerequisites

For image acquisition, a camera device is required whose image sensor has to be able to capture the whole dynamic range of the projector without clipping. If this precondition is satisfied, no further calibration is required. As with other camera based calibration algorithms, a color channel adjustment should be carried out to adjust the projector's white point to match the desired perceived impression of the human observer¹. It should be noted that, like in any camera based photometric projector compensation technique, the resulting quality largely depends on the camera used.

2.2. Color Mapping Method

Instead of assuming a linear color relationship described by a per-pixel matrix multiplication, we define it by a color mapping function:

$$c_c^* = f(c_{in}^*) \quad (1)$$

transforming the input colors c_{in}^* into the compensation colors c_c^* required to generate the desired intensities of c_{in}^* on the camera's image plane (* denotes all three *rgb* color channels). Our algorithm uses TPS interpolation to define

¹To accurately achieve desired color values independent of the camera settings used, a color calibration using a color checker chart can be applied beforehand, but this is not required for the algorithm to work.

this function. To gather the required parameters, the color space of the projector is sparsely sampled in a regular manner. While this regularity does not have to be strict, it should sample the extremes as well as the interior of the full *rgb* color cube evenly. This is achieved by generating colors with n increasing intensity levels from 0 to 255 in all three color channels as well as their combinations which leads to n^3 color samples. If knowledge about the rough shape of the projector response curve exists, the samples can be adjusted accordingly. For our algorithm, we used either $n = 4$ or 5 (cf. section 4 for an error analysis). These samples are sequentially projected and captured by the camera. After acquisition, the input colors and the corresponding captured images are used to compute the weights for each per-pixel interpolation function:

$$f(c_{in}^*) = \sum_{i=1}^N \omega_i^* \varphi(\|c_{in}^* - c_i^*\|) + \omega_{N+1}^* + \sum_{i=N+2}^{N+4} \omega_i^* \cdot \begin{cases} c_{in}^r, i = N + 2 \\ c_{in}^g, i = N + 3 \\ c_{in}^b, i = N + 4 \end{cases} \quad (2)$$

where c_{in}^* is the input color sample, N is the number of captured color samples, ω_i^* a series of $N + 4$ weighting coefficients per color channel, $\|\cdot\|$ the distance in Euclidean space, c_i^* the captured reference color samples, and φ is chosen to be the TPS RBF:

$$\varphi(d) = \begin{cases} 0, & d = 0 \\ d^2 \log d, & \text{otherwise} \end{cases} \quad (3)$$

which minimizes the integral of the squared second derivative of $f(c_{in}^*)$ and thus is well suited to generate a smooth color mapping. d is the Euclidean distance of the normalized sample color c_{in}^* to the normalized reference colors c_i^* . The per-pixel computation of ω_i^* has to be carried out only once per projector-camera setup and is calculated as described in Donato *et al.* [6]. To reduce the influence of outliers in the measurement resulting from noise and sampling, a small regularization term is added to the thin plate spline weights computation as proposed in the same paper. During run-time, equation 2 has to be evaluated once for each surface point to calculate the projector pixel intensities for the compensation image.

2.2.1 Photometric Compensation

Having computed the weighting factors for each pixel, arbitrary input images can be transformed from the camera's into the projector's color space by computing the color transformation using equation 2. While this, in theory, accurately approximates the desired color values, errors might occur because of out-of-gamut clipping which might arise




Figure 2. Another result of the proposed method: (a): input image; (e): highly non-uniform projection surface (illuminated white); (b): captured projection of (a) onto (e); (c): compensation without adjustments shows significant clipping errors. (d): reduced clipping errors after applying the global optimization step. (f-h): geometrically warped projection images generating camera images (b-d).

from too bright or low *rgb* values that cannot be reproduced on dark or colorful surface pigments or on grounds of ambient illumination. To avoid this, the input image can be globally adjusted by adapting the overall brightness and saturation such that the color values required for the compensation still can be generated by the projector. As already described in [2] an adaptive, spatially varying adjustment of the input image, however, has the potential to increase the overall perceived image quality even further, especially if the surface contains high spatially reflectance variations as the one shown in figures 1 and 2.

2.2.2 Global Optimization

Two main factors influence the occurrence of regional clipping errors: on the one hand they depend on the local reflection properties of the surface, while on the other hand the intensities of the image content are crucial. While the former is static, the latter can be reduced by optimizing the input colors and slightly changing the content similar to the approaches presented in [2, 9]. We smoothly adapt the luminance of the input colors by a spatially varying scaler p to avoid clipping. To achieve this goal, a non-linear optimization is applied to the input image minimizing the sum of the following per-pixel errors:

- Saturation error, occurring due to limited maximum projector brightness and will generate perceived image artifacts:

$$err_{sat}(x, y) = \begin{cases} (c_c^*(x, y) - 1.0)^2, & c_c^*(x, y) > 1.0 \\ 0, & else \end{cases} \quad (4)$$

- Intensity error, resulting from an intensity reduction at the current pixel which will reduce the image brightness (an increase in intensity is accepted):

$$err_{int}(x, y) = \begin{cases} (1.0 - p(x, y))^2, & p(x, y) < 1.0 \\ 0, & else \end{cases} \quad (5)$$

- Gradient variation error resulting from the spatially varying intensity adjustments leading to potentially visible local intensity variations:

$$err_{grad}(x, y) = (p(x, y) - p(x - 1, y))^2 + (p(x, y) - p(x + 1, y))^2 + (p(x, y) - p(x, y - 1))^2 + (p(x, y) - p(x, y + 1))^2 \quad (6)$$

Independent weights are applied to these errors to generate an acceptable tradeoff between the image degradation from clipping errors as well as global and smooth local luminance reduction in the final error term:

$$err = \sum_{x=0}^{width-1} \sum_{y=0}^{height-1} \omega_{sat} \cdot err_{sat}(x, y) + \omega_{grad} \cdot err_{grad}(x, y) + \omega_{int} \cdot err_{int}(x, y) \quad (7)$$

We used error weights of $w_{sat} = 200$, $w_{grad} = 1$, and $w_{int} = 50$ for our setup.

To speed up the optimization process, the optimization is only applied to a sub-sampled input image. This is applied such that only the darkest value in each image rectangle defining one pixel in the low-resolution representation is stored for the compensation data calculation, while only the brightest value is stored for the input image. This ensures that the worst case is considered during the optimization step. Currently we use a sub-sampling of 40×30 pixels which seemed to be a good trade-off in terms of accuracy as well as computation time for our test setups. Note that this sub-sampling should be adjusted depending on the spatial frequency of the surface texture. To solve this relatively large number of variables in a reasonable amount of time, a bound constrained optimization not requiring derivatives was applied [16]. After computation, the result is smoothly up-sampled into its original resolution and used to adjust the luminance of the image used as input for equation 2.

3. Experimental Setup

Different projector-camera pairs projecting onto multi-colored surfaces were used for evaluation. The ambient illumination was kept at a low and constant level. Results for

three different setups are shown in figures 1,2, and 5.

Hardware The method was tested using a LCD and a single-chip DLP projector ² with activated *brilliant color*[12] processing and a manually adjusted white point. For image acquisition, a Canon DSLR ³ was used. We captured JPEG images instead of RAW since our method does not require linearized images. Camera noise and color shifts from sequential DLP processing were minimized by averaging 8 images of each projection. For computation, a state-of-the-art PC was used ⁴

Geometric Calibration The mapping between projector and camera pixels was generated via a robust structured light scan based on gray codes with additional line shift patterns adapted from [10].

Implementation We implemented the algorithm in C++ with OpenMP based parallelization. For acceleration, the compensation was written using a GLSL fragment program as well. While it requires a large amount of GPU memory, it enables a significant speedup achieving real-time processing. Computation including image download took $\sim 40ms$ (compared to 7s in the CPU version) for a 1080p input image, which is sufficient for video playback.

4. Quality Comparison

We compared our method to the one presented in [3], which was applied by assuming the color spaces of the devices to be sRGB. While the camera used actually generates images encoded in this color profile, the DLP projector’s response curve was not described accurately by this profile. A comparison is given in figure 3. While the method mentioned also compensates for the majority of the errors occurring, it’s visibly apparent that the nonlinear method is able to more closely reproduce the desired colors if no accurate device calibration data is available.

²LCD: Epson TW3200, DLP: BenQ W1100

³Canon EOS 600D

⁴Intel i7 3930, 32GB RAM, Nvidia Geforce 670




Figure 3. Comparison to the method presented in [3] assuming a sRGB response of the devices. The absence of an accurate representation of the sRGB color space in the current settings of the DLP projector reduces the compensation quality in (a,d) compared to our approach (c,f). The input images are shown in (b) and (e). The surface of figures 1 and 2 has been used for the comparison.




Figure 4. Compensation accuracy with respect to the number of samples used to compute the TPS weights: The diagram shows the average ΔE_{00}^* [18] of the two sample image series shown on the left. While there is a clear quality improvement between 4^3 and 5^3 samples, the extension to 6^3 samples only marginally improves the result, while requiring almost twice as many images.

5. Summary and Conclusion

In this paper we showed that a high quality photometric compensation can be carried out without the need of any radiometric device pre-calibration. By projecting and capturing a reasonable number of images ⁵ a non-linear color mapping can be generated which enables an accurately compensated projection even on strongly textured surfaces. The proposed method offers several advantages compared to existing methods:

- Linearization errors resulting from inaccurate or noisy radiometric calibration don’t influence the compensation quality
- The algorithm is able to compensate for the complex, multi-primary processing of many single-chip DLP projectors
- No further calibration hardware is required which makes the compensation straightforward to deploy

The overall memory requirement, however, is significantly higher compared to linear mapping algorithms. In our system, we require up to $254 * 3$ floating point values for each pixel which, depending on the resolution, requires several gigabytes of memory. While this is not a serious problem on modern computers, the GPU implementation might suffer from this requirement on hardware with limited VRAM. We could, however, show that current consumer level GPUs are able to process this data in real-time. A comparison with other existing methods as well as a thorough analysis on the optimum tradeoff between input samples and compensation accuracy as presented in figure 4 is part of our further investigations. While the compensation is calculated in real-time, the global optimization takes, depending on the setup, five to 30 minutes to converge. An optimized GPGPU based

⁵currently we use up to 125 images for the data acquisition



Figure 5. Additional samples of the proposed method using difference setups. The upper row (a-d) shows the desired input images, the bottom row (e-h) the captured projections. In (i-k) the surfaces are shown under white projector illumination: In (a) a home cinema LCD projector is used, in (b-d) a DLP projector is used projecting from the right hand side.

non-linear optimization might even be able to achieve interactive frame rates including the content-dependent optimization. Adding local chrominance adaptation as well as content preservation in image regions darker than the ambient illumination to the optimization is another area of future research.

References

- [1] D. G. Aliaga, Y. H. Yeung, A. Law, B. Sajadi, and A. Majumder. Fast high-resolution appearance editing using superimposed projections. *ACM Trans. Graph.*, 31(2), 2012. **2**
- [2] M. Ashdown, T. Okabe, I. Sato, and Y. Sato. Robust content-dependent photometric projector compensation. In *Proc. of the 2006 Conference on Computer Vision and Pattern Recognition Workshop*. IEEE Computer Society, 2006. **2, 4**
- [3] O. Bimber, A. Emmerling, and T. Klemmer. Embedded entertainment with smart projectors. *IEEE Computer*, 38(1), 2005. **2, 5**
- [4] O. Bimber, D. Iwai, G. Wetzstein, and A. Grundhöfer. The Visual Computing of Projector-Camera Systems. In *Proc. Eurographics (State-of-the-Art Report)*, 2007. **1**
- [5] X. Chen, X. Yang, S. Xiao, and M. Li. Color mixing property of a projector-camera system. *ProjCams*. ACM, 2008. **2**
- [6] G. Donato and S. Belongie. Approximate thin plate spline mappings. In *Proc. of the 7th Europ. Conference on Computer Vision, ECCV '02*, 2002. **3**
- [7] J. Duchon. Splines minimizing rotation-invariant seminorms in Sobolev spaces. In *Constructive Theory of Functions of Several Variables*, volume 571 of *Lecture Notes in Mathematics*, chapter 7. 1977. **1, 3**
- [8] K. Fujii, M. Grossberg, and S. Nayar. A projector-camera system with real-time photometric adaptation for dynamic environments. 2005. **2**
- [9] A. Grundhöfer and O. Bimber. Real-time adaptive radiometric compensation. *IEEE Trans. on Visualization and Computer Graphics*, 14, 2008. **2, 4**
- [10] J. Guehring. Reliable 3d surface acquisition, registration and validation using statistical error models. In *3-D Digital Imaging and Modeling, 2001. Proc. Third International Conference on*. IEEE, 2001. **5**
- [11] R. L. Hardy. Theory and applications of the multiquadric-biharmonic method. 20 years of discovery 1968-1988. *Computers & Mathematics with Applications*, 19(8), 1990. **2**
- [12] D. C. Hutchison. Introducing brilliantcolor technology. Technical report, Texas Instruments, DLP Products. **5**
- [13] A. Majumder and R. Stevens. Color nonuniformity in projection-based displays: Analysis and solutions. *IEEE Trans. on Visualization and Computer Graphics*, 10(2), Mar. 2004. **1**
- [14] S. K. Nayar, H. Peri, M. D. Grossberg, and P. N. Belhumeur. A Projection System with Radiometric Compensation for Screen Imperfections. In *ProCams*, 2003. **1, 2**
- [15] S. H. Park, S. Yang, and B.-U. Lee. Adaptive chrominance correction for a projector considering image and screen color. In *Proc. of the 3rd international conference on Advances in visual computing, ISVC'07*, 2007. **2**
- [16] M. J. D. Powell. The BOBYQA algorithm for bound constrained optimization without derivatives. Aug. 2009. **4**
- [17] B. Sajadi, M. Lazarov, and A. Majumder. Adict: accurate direct and inverse color transformation. In *Proc. of the 11th Europ. conference on Computer vision, ECCV'10*, 2010. **3**
- [18] G. Sharma, W. Wu, and E. N. Dalal. The ciede2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research and Application*, 30(1), 2005. **5**
- [19] D. Shepard. A two-dimensional interpolation function for irregularly-spaced data, 1968. **2**
- [20] D. Wang, I. Sato, T. Okabe, and Y. Sato. Radiometric compensation in a projector-camera system based properties of human vision system. 2005. **2**
- [21] G. Wetzstein and O. Bimber. Radiometric Compensation through Inverse Light Transport. *Proc. of Pacific Graphics*, 2007. **2**
- [22] T. Yoshida, C. Horii, and K. Sato. A Virtual Color Reconstruction System for Real Heritage with Light Projection. In *Proc. of International Conference on Virtual Systems and Multimedia (VSMM)*, 2003. **2**