Radiometric Compensation of Nonlinear Projector Camera Systems by Modeling Human Visual Systems

by

Matthew Post

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Applied Science
in
Systems Design Engineering

Waterloo, Ontario, Canada, 2023

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Radiometric compensation is the process of adjusting the luminance and colour output of images on a display to compensate for non-uniformity of the display. In the case of projector-camera systems, this non-uniformity can be a product of both the light source and of the projection surface. Conventional radiometric compensation techniques have been demonstrated to compensate the output of a projector to appear correct to a camera, but a camera does not possess the colour sensitivity and response of a human. By correctly modelling the interaction between a projector stimulus and camera and human colour responses, radiometric compensation can be performed for a human tristimulus colour model rather than that of the camera. The result is a colour gamut which is seen to be correct for a human viewer but not necessarily the camera.

A novel radiometric compensation method for projector-camera systems and textured surfaces is introduced based on the human visual system (HVS) colour response. The proposed method for modelling human colour response can extend established compensation methods to produce colours which are human-perceived to be correct (egocentric modelling). As a result, this method performs radiometric compensation which is not only consistent and precise, but also produces images which are visually accurate to an external colour reference. Additionally, conventional radiometric compensation relies on a solution of a linear system for the colour response of each pixel in an image, but this is insufficient for modelling systems containing a nonlinear projector or camera. In the proposed method, nonlinear projector output or camera response has been modelled in a separable fashion to allow for the linear system solution for the human visual space to be applied to nonlinear projector-camera systems.

The performance of the system is evaluated by comparison with conventional solutions in terms of computational speed, memory requirements, and accuracy of the colour compensation. Studies include the qualitative and quantitative assessment of the proposed compensation method on a variety of adverse surfaces, with varying colour and specularity which demonstrate the colour accuracy of the proposed method. By using a spectroradiometer outside of the calibration loop, this method is shown to produce generally the lowest average radiometric compensation error when compared to compensation performed using only the response of a camera, demonstrated through quantitative analysis of compensated colours, and supported by qualitative results.

Acknowledgements

I would like to thank all those who made this research and the completion of my degree possible. The support received from my supervisor Paul Fieguth, whose patience, guidance and ceaseless dedication to excellence inspired me to become a better researcher, student, and teacher.

I would like to thank Professors John Zelek and Bryan Tripp for thaking time from their busy schedules to review my thesis.

I would also like to also thank Christian Scharfenberger and David Clausi, who first introduced me to the cutting-edge work of Waterloo's Vision and Image Processing Lab. Their encouragement and passion during my third and fourth co-op terms as an undergrad student steered me onto my current course of robots and computer vision work.

For their tireless support of my projector-camera work, both as a student and at Christie Digital, I must thank Mohamed Naiel, Mark Lamm, Zohreh Azimifar and Saed Moradi, particularly with their help with this thesis having good figures that tell the story.

Dedication

This thesis is dedicated to those who believed in me and never let me give up, my family and my friends.

Table of Contents

List of Tables

List of Figures

Nomenclature

$\mathbf{C}$	A vector representing the channel-wise camera colour response for a pixel, typically red, green, and blue
$\mathbf{e}(\lambda)$	The spectrum of light output by a particular projector colour channel as a function of wavelength, with the pixel intensity as a factor.
$\mathbf{F}(\cdot)$	An arbitrary nonlinearity function for a three dimensional input vector
$\mathbf{K}$	A matrix representing the compensation solution to be applied to achieve a desired projector output
$\mathbf{K_a}$	A matrix representing the compensation solution to be applied to achieve a desired projector output which is augmented to account for ambient light
$\mathbf{q}(\lambda)$	The spectral sensitivity of a given sensor, here either a camera colour channel or human tristimulus curve
$\mathbf{P}$	A vector representing the channel-wise projector colour output for a pixel, typically red, green, and blue
$\mathbf{s}(\lambda)$	The spectral reflectance of the surface upon which light is projected
$\mathbf{S}$	A vector representing the channel-wise colour response for a pixel in a standardized colourspace, typically red, green, and blue
$\mathbf{S_a}$	A vector representing the channel-wise colour response for a pixel in a standardized colourspace, typically red, green, and blue, augmented by appending a 1
$\mathbf{V_{c,p}}$	A matrix representing the mapping from camera colourspace to projector colourspace
$\mathbf{V_{s,h}}$	A matrix representing the mapping from the XYZ colourspace to a standard RGB colourspace

$\mathbf{V_{h,p}}$	A matrix representing the mapping from the projector colourspace to the XYZ colourspace
$\mathbf{w}(\lambda)$	The spectrum of light output by a particular projector colour channel as a function of wavelength, at maximum intensity.
λ	Wavelengths of light, typically in nanometers.

Chapter 1

Introduction

In recent years, the use of projection technology to display media on unconventional surfaces has become increasingly prevalent. Projectors are used in not only cinemas, but in planetarium domes, amusement park rides, and to project onto architecture. Projectors have also become cheaper and increasingly powerful, with some individual projectors capable of outputting 60,000 lumens.

With the increased availability of projectors with high luminous outputs, there has been an increased demand to project onto novel and challenging surfaces. Often these projection scenarios include an array of projectors illuminating a pre-existing structure or a surface not originally designed for projection. Examples are varied, but typically the surface is a building, monument, and can include a range of 3D geometry, textured surfaces, and image content specifically designed for the scene. These challenging scenarios each result in colours or textures in the projected imagery which are not part of the intended content. Such textures and colours can easily reduce the impact of the content on an audience. Similarly, blended multi-projector displays present another challenge to seamless content, which is projector colour matching. Two projectors, even projectors from the same brand and model, can present different colour gamuts due to differing light sources. This can produce visible seams in content, and make the content colours look incorrect with respect to that of other projectors. This typically requires the painstaking manual tuning of colours, which can be prohibitive for large arrays of projectors.

A notable example of projection on the 3-dimensional surface of a monument is the Statue of Unity in Gujarat, India. The 182 metre tall, bronze-clad statue was illuminated with 30 Christie Crimson projectors, each outputting an HD image to cover a portion of the surface. The surface poses a particular challenge to project on, as the bronze cladding

creates a surface which changes the proportion of light reflected based on relative angle. The result of projecting on this surface is content which lacks the intended colours, luminance of the surface varying widely, and textures of the surface visible through the projected content. The challenges of projection mapping can be seen in Figure 1.1b, showing projection on a flat yet challenging ice surface, and projection on the Statue of Unity in uncontrolled conditions. Both show different conditions in which non-uniformity of projection can be perceived by a viewer.

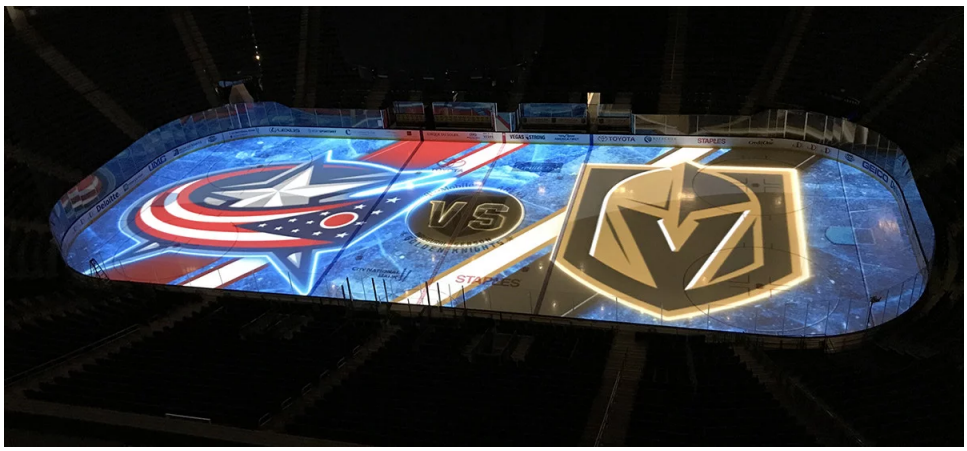

(a) Projection mapping on a 2 dimensional surface as demonstrated on the ice surface of a hockey rink.

(b) Projection mapping on a 3 dimensional surface as demonstrated on the Statue of Unity.

The statue presents obvious challenges to projection, not just for geometric alignment of pixels, but also for the uniform illumination. The geometry of the statue surface causes the angle of incident of the projector, and thus the effective luminance to change, and the

texture and colour of the surface causes further non-uniformity. The ice rink poses similar challenges to the statue due to differing angles of projection of the projectors, but has the additional difficulty of being a partially specular surface. This means that much of the projected light for a pixel is evenly scattered, the relative viewing angle can significantly change perceived luminance. Additionally, on a flat white surface, humans can more readily perceive differences in the colour between projector light sources, thus requiring the colour range of the projectors to be more precisely matched.

1.1 Thesis Contribution and Proposed Method

The target of this thesis is to propose a method for radiometric compensation which overcomes several shortcomings of existing solutions in three separate parts of the compensation pipeline. There are three major contributions:

1. Radiometric compensation of non-linear projector-camera systems is solved efficiently. This method is an improvement over existing methods by virtue of separating the projector non-linearity and linear radiometric compensation. This method consists of capturing the range of a projector's colour gamut using a camera, and representing the recorded values in terms of the projector primary colours. This method is detailed in Chapter 3.

2. The process of colour correction of projectors with respect to a human's sensitivity is introduced. Projector systems are compensated by modelling the differences in interaction between the way projector light sources interact with cameras and the human visual system. The colourspace transformations between the colours of cameras and projectors calibrate projectors to a common gamut. This method is outlined in Chapter 4.

3. The detection of pixel-level outliers in radiometric compensation is introduced. Both the detection and the correction of these pixels is explored. This method identifies poorly compensated pixels for a given projector, and uses information from nearby pixels to fill in the missing information. This method is outlined in Chapter .

The contributions in this thesis are each proposed to solve a particular problem in a radiometric compensation pipeline, with the goal of being eventually used in commercial applications. The solutions proposed must result in an comprehensive pipeline for radiometric compensation which can be applied to high-resolution and high-framerate projectors

in realtime, and can be run without human intervention and tuning. The performance of the proposed methods is demonstrated in Chapter 6.

Chapter 2

Background

In this chapter, background information and related work in this thesis is presented. The principles of ubiquitous projection, projection on any surface, and projector-camera systems are explored. Section 2.1 provides an overview of existing compensation methods and the underlying mathematics for projector camera systems. Section 2.2 describes the process of determining the geometric relationships between projector pixels and camera pixels. Section 2.3 describes the spectral stimulation of the human visual system, and the principles of the projector spectral outputs. Existing work in the field is described, and shortcomings are assessed. Section 2.4 provides an overview of existing compensation methods and their deficits.

2.1 Overview

Radiometric compensation is the process of altering the output of images on a display to compensate for non-uniformity of the display. A display may take the form of a computer monitor, an LED array wall, a video projector, or another visual medium. In the case of video projectors, this non-uniformity can be a product of non-uniformity of both the light source and of the projection surface. Radiometric compensation aims to achieve a more accurate representation of the projected image by modelling non-uniformity. In projector-camera systems, the camera provides the feedback required to allow the projector image to be compensated for background texture of a projection surface. In the vast majority of these systems, cameras and projectors use a three-channel RGB representation of images.

In the literature, several methods for radiometric compensation have been proposed and have been shown to work effectively in real-time with projector-camera systems with linear

relationships [21, 2, 23]. In general, such methods work by determining a multi-channel linear model of the reflected light for each pixel and determining the inverse solution. To create a closed calibration loop for each pixel on a given projection surface, pixel correspondences are first calculated between the projector and camera through the use of structured light in a set of training images. Next, colour training images are projected on the surface and captured with the camera. These training images are then used to constrain the solution to the system inverse model for each pixel. The result is a 3×3 matrix model V for each pixel, where each element relates the 3 spectral contributions of the projector to each of the 3 colour responses of the camera. The matrix would be a diagonal matrix if each projector colour channel only stimulated the respective camera colour channel. The projector and camera colour channels both have overlapping distributions, necessitating the off-diagonal terms in the matrix V. Let C and P denote a 3×1 vector of RGB values output from the camera and input of the projector, respectively. The inverse solution to a given system provides the required projector output, P, from a desired camera colour as follows:

$$C = VP, \quad P = V^{-1}C \tag{2.1}$$

This model of a pixel's colour in a projector-camera system does not include parameters for ambient light, and thus can only model systems where the projector is the only source of light for the camera. This was improved by Yoshida [31] with the augmentation of the matrix equation with a constant term to include ambient light contributions as follows:

$$P = KC_a, \quad C_a = \begin{bmatrix} c_r \\ c_g \\ c_b \\ 1 \end{bmatrix} \tag{2.2}$$

where K is a 3×4 matrix that represents the inverse of the pixel colour response. Alternatively, certain projectors have nonlinear responses to inputs which prevent simple linear system modeling. The nonlinearity of the projector output causes the radiometric response of each projected pixel to become nonlinear. This nonlinearity is designed into many projectors output response to enhance colour vibrancy and perceived contrast, as with a gamma correction applied to other displays. Nonlinearity can include gamma corrections, white boosting, colour boosting, or S-curves applied to the colour intensities, all which may be desirable for some enhancement of the projected image, but less desirable for those who might be trying to model the system. Recently, Grundhöfer and Iwai [6] proposed a method for radiometric compensation of such non-linear systems using a lookup table per pixel and thin plate spline interpolation to approximate the colour inverse problem. The lookup table inherently provides solutions for nonlinear projector response

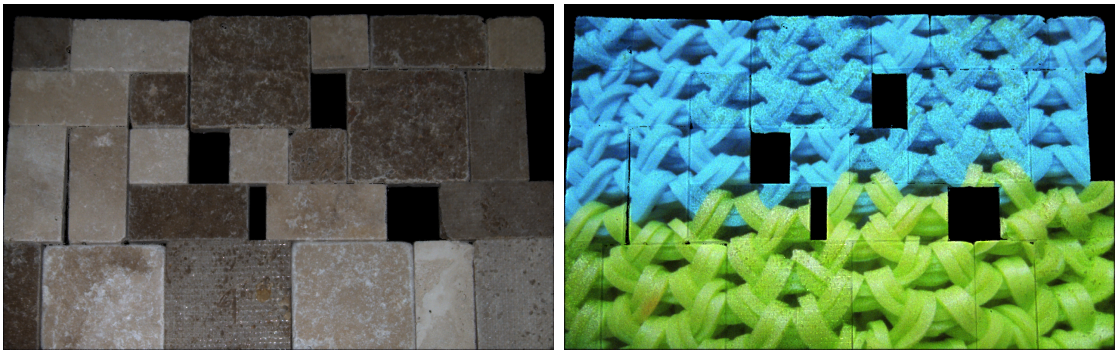

(a) (b)

Figure 2.1: (a) Sample *Stone Tile* test background and (b) compensated image using the proposed scheme. This experiment is conducted in high ambient light conditions.

and ambient light. However, this method has two significant drawbacks; firstly, the sparse lookup table of each pixel's colour response must be stored, and secondly, a spline interpolation must be performed for each pixel in each image to be compensated. This requires a significant increase in both computations and storage per pixel over the linear solution.

Here, a method that efficiently computes the radiometric compensation of a nonlinear projector system is proposed. This method separates the nonlinearity of the projector from the linear radiometric compensation. The proposed scheme provides a significant reduction in computational complexity and storage requirements of the compensation when compared to that of the method in [6]. This reduction is achieved without sacrificing the accuracy of the compensation; a sample of this nonlinear compensation can be seen in Figure 5.1, where a stone tile wall is being illuminated with high ambient light.

2.2 Geometric Calibration

Geometric calibration of projector-camera systems is required in order to establish the relationships between their respective pixels. In the context of radiometric compensation, this relationship is required to use the information gathered with the camera to compensate specific projector pixels. The geometric relationship can be expressed in multiple ways including correspondences between camera pixels and projector pixels or the pose of the devices relative to each other when the screen geometry is known. The geometric calibration generates a mapping between camera pixels and projector pixels, allowing the

camera to be used for calibrating the projector colours in a secondary calibration procedure. The projector-camera geometric relationship is calculated assuming a fixed scene, in which the relationship between pixels of both devices does not change. The correspondences between pixels are sometimes sufficient to accomplish all calibration tasks, such as radiometric compensation. In the case that specific video content must be aligned to the surface geometry, the relative poses between projectors and cameras is calculated from the pixel correspondences. In either case, the pixel correspondences typically need to be as dense as possible, and in the case of radiometric compensation, must be pixel-dense in the projector domain.

Geometric calibration is often accomplished through the use of structured light. Structured light is any pattern with known geometric properties that is projected and allows pixels in a camera image to be related back to the projected pixels. Numerous methods of encoding spatial information in projected light have been developed [27] and demonstrated to be effective in projector-camera systems where the geometric relationships must be known [2, 8]. Such methods can include Gray Coded Binary patterns (GCB), blob patterns, or sinusoidal waves [27]. In these cases, the intensity of the pixels (on or off for binary) allow each camera pixel to be decoded to a projector pixel given a temporal encoding. The projector is used to display a sequence of images, and the sequence of intensities associated with a given pixel can be used to identify the pixel uniquely in a camera image. By covering the entire projector field with temporally-encoded pixel addresses, pixel-dense mappings between projector pixels and camera pixels can be obtained.

Though the design and development of structured light test patterns falls beyond the purview of the work presented, it is a prerequisite for the work. In the case of Gray Coded Binary patterns (GCB), the horizontal positions and vertical positions of pixels are encoded separately, with two sequences of test patterns. In each sequence, pixels are switched to either fully on (white) or off (black), based on Gray (reflected binary) code. Gray coding has the property that any two adjacent binary codes in the set differ by a single pixel value. By direct consequence, a single bit error in decoding means a single pixel error in decoded position. The encoding method forms bar patterns as seen in Figure 2.2. Each of the columns show a different sequence of horizontal bars used to encode vertical pixel position in the projected images.

The test patterns used for generating the correspondences between camera and projector must be selected based on the requirements of the calibration. Projection on a 3-dimensional, discontinuous surface requires a pixel-dense test pattern sequence such as the Gray-coded binary patterns outlined above, though this requires the use of a longer test pattern sequence. Projection on smooth, continuous surfaces does not require pixel-dense correspondences, and thus shorter sequences, or patterns which rely on spatial continuity

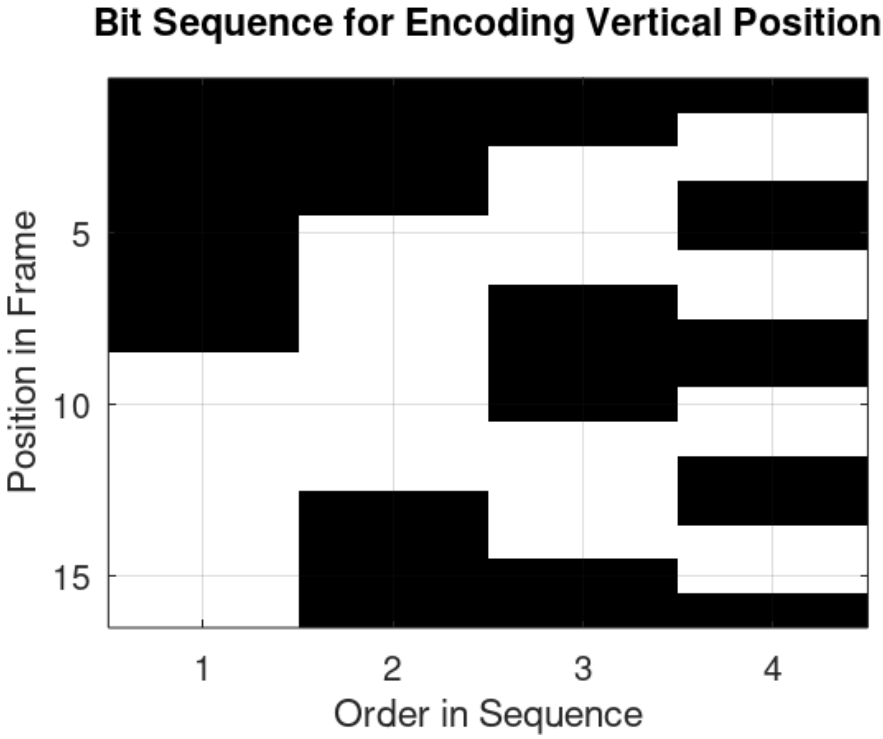

Figure 2.2: Sample of Gray-Coded binary patterns for 16 different vertical positions, with 4 different horizontal bar patterns.

of the pattern in the camera domain may be used. For the purpose of radiometric compensation, pixel-dense correspondences between projector and camera are required such that each projector pixel may be compensated appropriately for a viewer.

2.3 Human Visual System

The human eye is a remarkable difference detector, able to quickly discern edges in luminance and colour in a scene, whether this change occurs spatially or through time. Radiometric compensation aims to deceive this detector into seeing spatial uniformity by producing colours on a surface which are perceived to be the same, while possibly having different light spectra. The correct compensation of an image not only produces a uniform image in a scene, but also an image of the correct colours. The uniformity indicates accuracy, the specific colour indicates precision. A radiometric compensation technique must

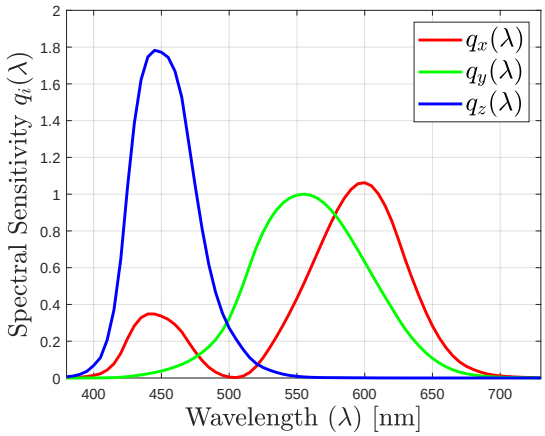

Figure 2.3: Human spectral sensitivity. This is the human visual system tristimulus model which became the CIE 1931 standard model. This model is still prevalent in industry, despite not being fully representative of the human population.

thus correctly account for the different projection spectra and the spectral sensitivity of the human eye.

The spectral sensitivity of the human visual system is typically modelled as a tristimulus system, composed of three response functions each of which approximates the sensitivity of a particular set of cones in the human retina. These curves are traditionally represented as $X(\lambda)$, $Y(\lambda)$, $Z(\lambda)$. To maintain generality of the equations of spectral responses, these will be instead be represented by $q_X(\lambda)$, $q_Y(\lambda)$ and $q_Z(\lambda)$ respectively. These spectral response curves can be seen in Figure 2.3. The three curves roughly cover the wavelengths which humans perceive as red, green, and blue respectively. Note the two maxima of the $q_X(\lambda)$ in figure 2.3, which causes perceived colours to wrap from red to blue through magenta, showing the human response of $q_X(\lambda)$ to both short and long wavelengths, corresponding to blue and red respectively.

These curves can be combined with the spectra of incoming light $r(\lambda)$, to approximate the colours which a human would perceive using an inner product:

$$C_i = \int q_i(\lambda) r(\lambda) d\lambda \tag{2.3}$$

The result of this inner product of the spectrum of incoming light with the spectral response of the human visual system is dependant on the shapes of the curve for both.

This means that the perceived colour computed from a measurement of the light spectrum will vary with the HVS model used. The way the human eye perceives colour cannot be directly measured, but rather it must be inferred statistically by human test subjects matching perceived colours and intensities of calibrated light sources [12]. Due to the nature of this statistical data collection, multiple human perceptive colour models have been developed [12, 2], each being demonstrated to be a suitable model for a particular test population. The differences in these models however lead to the calculations of colour perception not being consistent between the models, and so any method based on these models will vary in performance accordingly. The compensation technique developed in Section 4 relies heavily on a pre-existing model, and the outcome of compensation can be affected by the characteristics of the selected model.

2.4 Current Radiometric Compensation Techniques

The industry of projector displays extensively utilizes light projection on a wide variety of surfaces to paint a virtual canvas [20, 3, 16, 29]. Projector displays have a wide range of applications including cinema displays, amusement rides, or to project video content onto arbitrary surfaces that can range from the flat sides of buildings, to cars, museum interiors, sports arenas, and abstract sculptures. When these surfaces have colourations or textures, the effect of image projection onto these surfaces is clearly influenced by the background texture [23]. Projector-camera systems [23, 6] are used to provide feedback to compensate and adjust the projected image to reduce or, ideally, eliminate the effects of the background.

Conventional projector radiometric compensation schemes [6, 31, 2] are limited by the colour sensitivity of the calibration camera, since the subjective balance of the camera colour components determines the actual colours which are produced on-screen, resulting in compensation which appears incorrect to a human viewer. Thus, the addition of a mapping from camera tristimulus colour to human tristimulus colour will allow radiometric compensation to be performed with respect to the colour sensitivity of a human viewer rather than that of a camera. To the best of our knowledge, no method for radiometric compensation has been presented which can calibrate projector colours for a human viewer rather than that of a camera.

In this thesis, a novel projector radiometric compensation method is proposed that utilizes a mapping between the spectral sensitivities of a camera and a human viewer. This method introduces the mapping of a camera colour response to human colour response using

a projector stimulus for the first time. This radiometric compensation method is developed for linear systems and subsequently extended to nonlinear ones.

2.5 Related Work

The purpose of performing radiometric compensation with a projector-camera system is to provide perceived uniformity in images seen on the projection surface. This can be done either to compensate for non-uniformity of the projection surface [30] or of the light source itself [6]. This compensation is generally accomplished by using the camera in the system to determine the relationships between commanded projector intensities and colours observed on the surface.

Early methods for radiometric compensation were proposed using per-pixel linear solutions of increasing complexity [2, 31, 23, 30]. For instance, in [30] a method for compensating single-channel systems is introduced, which solves for a single gain value to be applied per-pixel in each target image, thereby compensating for the luminance but not compensating multiple colour channels simultaneously. Others proposed three-channel colour compensation [23, 31] by including colour mixing between projector and camera. These methods aim to produce colours to appear correct to a camera by modeling the light contributions from each projector channel to each camera channel by use of a 3×3 matrix solution per pixel. Ambient light was incorporated into radiometric compensation methods through the introduction of an additional column to the compensation matrices [31]. Recently, several methods have been proposed for radiometric compensation of nonlinear systems [16, 6, 25]. For example, a method for making the solution for per-pixel nonlinear systems feasible has been introduced in [6], where a lookup table is used to provide the compensation for each pixel. Later in [25], a method was proposed for separating the nonlinearity of the system from the pixelwise colour response to greatly reduce both the memory and computational requirements imposed by the nonlinear solution in [6]. These methods [6, 25] linearize a projector-camera system with a nonlinear response, but do not consider this for the nonlinearity or colour response of the human visual system.

Recent compensation methods [30, 8] have attempted to provide better compensation for viewers by modelling the human visual system. These methods improve visual quality of a scene by reducing edge artifacts from backgrounds in images, placing focus on contrast and edge sensitivity, but do not model human colour perception. For instance, the method outlined in [30] reduces edge artifacts due to background by dynamically scaling the intensity of a target image until the perceived edges are reduced in strength to below the threshold of human perception. However, this method does not perform radiometric

correction to preset target colours, and instead aims to make the camera colour response equal to the colour of the target image while reducing luminance and contrast to reduce edge artifacts from the background. The effects of perception to colour were outlined in [1], where compensation of content is performed in a perceptually uniform space. This method does provide a means of adapting the compensation based on the content of the image, but does not account for the human perception of the projector colours.

The problem of luminance and colour non-uniformity was explored in [18], but the proposed solution did not maintain appropriate white balance in compensated images when the projector colour channels were solved independently and instead perform single-channel luminance balancing. Most of the existing radiometric compensation and screen colour correction techniques [19, 11, 6, 29] have optimized the compensated projector outputs to appear accurate to a white point of a camera, or used a previously colour-calibrated projector. All of the methods mentioned above have not yet been extended to include the human visual colour response, and thus can be improved by better modeling of the human visual system.

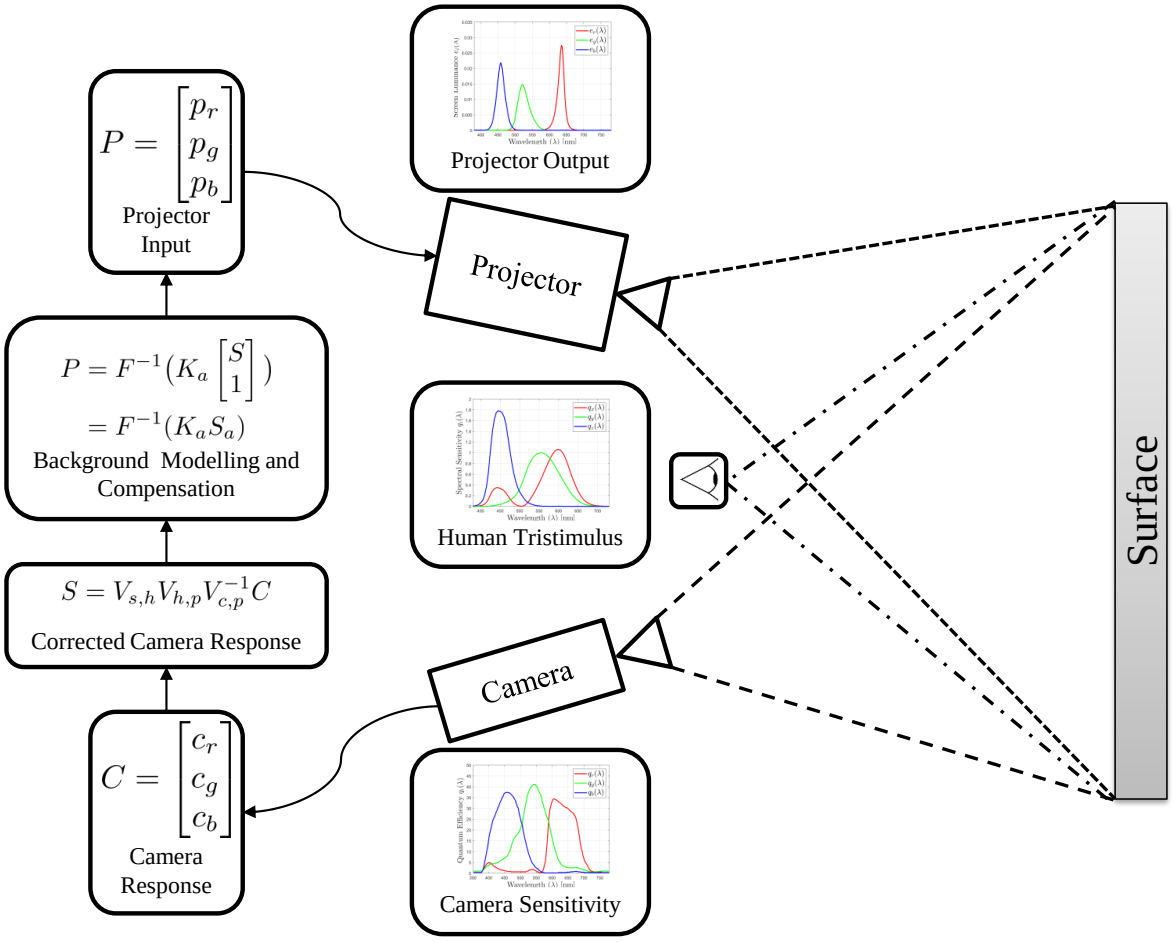

Figure 2.4: An overview of the proposed radiometric compensation scheme: Although a camera is used as feedback, the compensation scheme *explicitly* targets the human visual space, and not the camera space.

Chapter 3

Nonlinear Projector Modeling

Modern projectors have been optimized to provide the perceptually best video experience by increasing resolution, increasing framerate, and by modifying their projected colours. This modification of colours is typically done to increase the contrast of the projected image, in addition to the saturation of colours. This sort of modification involves the introduction of a non-linearity to the colour intensities of the projected image, such as a gamma function or a logistic function. Both increase contrast by reducing the intensity of dark colours towards black, and increasing the lighter colours towards white. Often this is termed *white boosting* or *colour boosting*. This nonlinearity introduces difficulty into modelling the system, as it invalidates linear matrix solutions for radiometric compensation.

Modern cinema class projectors such as the Christie Boxer and D4KLH60 operate at a resolution of 4K (4096×2160) and frame rates of 120Hz. These projectors require more computationally efficient nonlinear compensation methods than those that previously existed [6, 5]. In this chapter, a nonlinear radiometric compensation method capable of being applied to these high-performance cinema projectors is introduced. The proposed method uses a single lookup table to linearize the projector response, rather than using a lookup table per pixel as in [6]. The proposed method then employs the linear formulation proposed by Yoshida [31] which uses 3×4 matrices to model the linear radiometric response of a camera to both a projector and constant ambient light sources. The linearization step in the proposed scheme therefore does not need to be calculated for each pixel, but rather for the light source and then applied to the linear response model of all pixels.

3.1 Background Modeling

It is assumed that the the screen can be modelled by Lambertian reflectance [13]. This model holds valid for projection surfaces which reflect light in a more diffuse way than specular, which includes most conventionally designed projection surfaces. This assumption is not valid for specularly reflective surfaces, such as metal, glossy paint or glass, but these surfaces do not lend themselves to projection as the projected images are not well-formed on the surface. Thus when assuming lambertian reflectance, the screen irradiance for a given projector primary channel is a product of both the given pixel intensity for that channel and the spectral response of the same channel. The nonlinearity of the projector output can be caused by additional colour channels in the projector which do not directly correspond to RGB, such as cyan, magenta or white. This can also be caused by nonlinear gain applied to the colour channels such as a gamma function. Given the wavelength distribution $w_i(\lambda)$ of each colours channel i, the combined nonlinear irradiance e_i can be expressed as:

$$e_i = F_i(P(\overline{u})) \cdot w_i(\lambda) \tag{3.1}$$

Here, $\overline{u}$ is the pixel indices (u, v) in the projector frame of reference, and $F_i(\cdot)$ is the nonlinear projector response to input $P(\overline{u})$.

The background compensation for the projector is calibrated based on the colour feedback from the camera in the system. When the projector stimuli and camera responses can be treated as linear with respect to one another, the system and its inverse can be modelled as shown in (2.2). The colour and texture of the background surface for each pixel provide a simple scalar gain to each colour channel component, and so the system with background colouration and the corresponding compensation (the system inverse) can also be modelled as such.

This linear solution for radiometric compensation can thus be applied to a non-linear system if the correct linearization of the projector is known. Let the response of the camera be represented as C, the projector input as P, and the projector's internal nonlinearity function as $F(\cdot)$. The relationship of the camera response to the light reflected off of a projection surface can be described by the 3×3 matrix relationship V:

$$C = VF(P) \tag{3.2}$$

Then, the compensation of the system can be accomplished by inverting the system to solve for the inverse relationship K, allowing the required projector colour intensities P given an arbitrary desired camera response, C:

$$P = F^{-1}(KC) \tag{3.3}$$

Equation 3.3 can then be expanded to account for scenarios with ambient light on the projection surface by using an augmented camera response C_a as found in Equation 2.2, resulting in the nonlinear compensation:

$$\boxed{P = F^{-1}(KC_a)} \tag{3.4}$$

This equation is fundamental for fast nonlinear radiometric compensation; the linear per-pixel compensation matrix K is separable from the projector linearization function $F^{-1}(\cdot)$. The linear compensation method in (2.2) can then be used when a suitable model of $F^{-1}(\cdot)$ is known, which can in fact be learned.

The full system is then solved for each pixel by using an initial set of eight projected flat colour test patterns using the pseudo-inverse. The test patterns consist of the extreme projection values of the projector, namely white, the projector primaries and secondaries, and black. In the case of most projectors, the primary light drives are red, green and blue. The resulting secondary colours are cyan, yellow and magenta. These eight colour values are the eight vertices of the colour gamut of the projector, and every other possible colour output lies within the resultant cubic colourspace. In the case of solving the linear system for each pixel, more test colours provides greater stability to the solution of K. The vertices of the projector's gamut were selected to improve solution stability, but the system can be solved with just the projector's primaries and a black level image, as white, and the projector secondaries are linear combinations of red, green and blue.

3.2 Projector Modeling

The nonlinearity function $F(P)$ of a given projector is modelled through the use of a single lookup table to linearize the colour response for the projector. A lookup table was selected as it provides low constant-time access to all values for the function model, while using very little memory. A fully dense lookup table for every possible 8-bit RGB colour value only requires 3-byte colour values for each of the $(2^8)^3$ possibilities which is $3 \times (256)^3$ bytes, or 48 megabytes; such a data structure is small enough to fit in most modern GPU memory without concern. The lookup table of the proposed scheme is generated by sparsely sampling the projector's output over the full range of the projector's RGB channel inputs as follows. Single colour images are first projected on a near white surface, and captured with the camera. The camera images are then averaged over a 200×200 pixel patch in the center of the projected field. This provided a large enough sample size once averaged to produce a stable colour value even with significant camera sensor noise. Due to the

differences of the response of the camera with respect to the stimulus of the projector, a meaningful inverse of the system must provide the projector response only in terms of projector RGB input. The solution of the projector inverse model should thus be invariant to the camera response. To achieve this, all of the sample colours as seen by the camera are remapped with respect to the projector primaries, in this case the maximum outputs of red, green and blue channels. This mapping is computed by taking the inverse of the appended RGB camera responses V_0 to projector primaries as follows:

$$P_{out} = V_0^{-1} C \tag{3.5}$$
$$= V_0^{-1} V_0 F(P_{in}) \tag{3.6}$$
$$= F(P_{in}) \tag{3.7}$$

$$\text{where} \quad V_0 = \begin{bmatrix} v_{RR} & v_{RG} & v_{RB} \\ v_{GR} & v_{GG} & v_{GB} \\ v_{BR} & v_{BG} & v_{BB} \end{bmatrix} \tag{3.8}$$

Here, P_{out} is the nonlinear projector RGB output, C and P_{in} are the camera response and the projector RGB input in (3.2), and v_{ij} is the response of camera channel i to stimulus of projector channel j at maximum intensity. As shown in equations (3.5)-(3.7), the projector nonlinearity function $F(\cdot)$ can be determined by evaluating (3.5) to obtain the normalized output P_{out} for a given input P_{in}. The matrix in (3.8) is then used to map all colours seen by the camera to a colourspace of which the projector primaries are basis vectors.

The mapped colours are thus projector outputs P_{out} in terms of projector primaries, with known input RGB vectors P. This known relationship between inputs and outputs of the projector allows a lookup table to be built with no further dependence on the camera response but is solely a product of the projector configuration. This lookup table can be built from either dense or sparse sampling, with a suitable interpolation being performed prior to radiometric compensation. In the case of sparse sampling, the projector gamut can be sampled at a subset of the possible RGB inputs, and interpolated to derive the desired lookup table. In this thesis, linear interpolation with sample points chosen on a uniform 3D grid is used, resulting in n^3 points.

Background Compensation: To project a given target image on a given surface, radiometric compensation must be applied to each pixel. The compensation for a given projector and projection surface is performed by computing (3.3) for each desired pixel colour C in the input image. From (3.5) the function $F^{-1}(\cdot)$ which linearized the system

is approximated via a lookup table $\hat{F}(P_{out})$, and is computed once per projector configuration. The pixelwise linear compensation K can be obtained from (2.2) once for each projector-background setup. The evaluation of the proposed compensation method in (3.3) requires only a $(3 \times 4) * (4 \times 1)$ matrix multiplication and a single indexing operation per pixel in each projected frame. Comparatively, the existing state-of-the-art method in [6] requires a thin plate spline interpolation of 6^3 points per pixel.

Chapter 4

Projector to Human Colour Response Modeling

Modelling the colour response of the human visual system provides an alternative to a camera as a sensing method for performing radiometric compensation. Figure 2.4 shows an overview of the proposed scheme, consisting of a projector, camera, human viewer, computer and a surface to be compensated. In the proposed method, the colour response of a camera is corrected to accurately represent the colours as seen by the human viewer. The human colour response is subsequently mapped to the colourspace of the desired projector content, allowing the radiometric compensation for a given background and projector to be performed with respect to the human eye, and not just to a camera. The proposed radiometric compensation scheme will be developed to determine an independent solution for each individual pixel on a given projection surface.

4.1 Representation of Spectral Responses

In this section, the relationships that govern the spectral responses for the human eye and a camera with respect to a projector stimulus are presented. The human colour response to light stimuli can be represented as a tristimulus model [12] consisting of three gain functions as seen in Figure 4.1.

Let the spectral output of a projector $e_j(\lambda)$ be defined as

$$e_j(\lambda) = P_j w_j(\lambda), \tag{4.1}$$

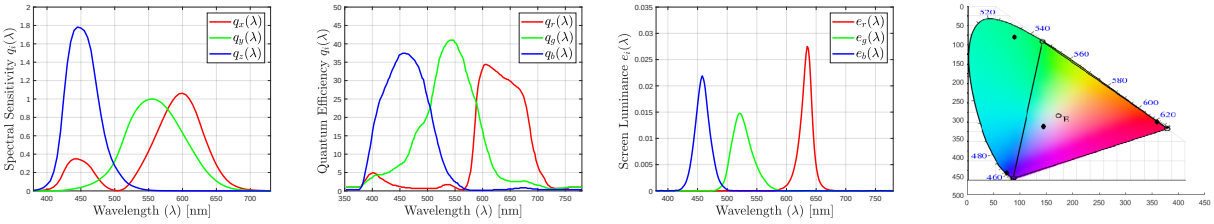

(a) (b) (c) (d)

Figure 4.1: An illustration of example spectral responses for human tristimulus curves (a), camera tristimulus curves (b), and the spectral output of a Christie Matrix StIM LED projector (c). Further, (d) shows the projector primaries in xy colourspace obtained for (c) as seen by (a).

where $w_j(\lambda)$ is the light intensity output of the jth stimulus (such as a projector colour channel) as a function of wavelength λ, and P_j is the proportional gain (pixel intensity) corresponding to this jth stimulus. The proportional gain P_j is the jth element of the RGB input to the projector, P, used previously. Next, let $q_i(\lambda)$ denote the response of the ith sensor (such as a camera colour channel or human retinal cell type) as a function of wavelength λ. It has been shown [23] that the resultant response, r_i, of the ith sensor to the combined reflected stimuli can be expressed as

$$r_i = \sum_j v_{i,j} P_j \tag{4.2}$$

where

$$v_{i,j} = \int_{\lambda_{\min}}^{\lambda_{\max}} q_i(\lambda) w_j(\lambda) s(\lambda)\, d\lambda. \tag{4.3}$$

Here $s(\lambda)$ is the surface reflectance for a given light wavelength, and $\lambda_{\min}$ and $\lambda_{\max}$ are the bounds of the sensor wavelength sensitivity. Examples of the response curves $q_i(\lambda)$ and projector stimuli $e_j(\lambda)$ can be seen in Figure 4.1 (a-c). It must be pointed out that both the camera response (quantum efficiency) and projector stimuli (spectral output) are highly dependent on the characteristics of the respective devices.

As derived in (4.3) the sensor tristimulus colour response, $v_{i,j}$, is represented by the inner product of three colour sensitivity functions $q_i(\lambda)$ of a given camera or human colour perception model (such as CIE 1931 [12] and CIE 1964 [24]) and the spectral outputs of each projector channel $w_j(\lambda)$ with a given screen spectral reflectance $s(\lambda)$. The projector spectral output $w_j(\lambda)$ can not be measured directly, due to the sensitivity of most sensors

and the power of the projectors to be modelled. However, from (4.1) it follows that

$$w_j(\lambda) \propto e_j(\lambda) \tag{4.4}$$

where $e_j(\lambda)$ can be approximated by the measurements of a spectroradiometer. The projector spectra only need to be measured once and stored for each projector type, which removes the requirement of measuring the spectrum at the time of radiometric compensation.

Let C and H be the spectral responses of the camera sensor and human visual system, respectively, and let P be the channelwise input to the projector. For instance, these vectors can be expressed in terms of three primary channels, RGB or XYZ, as

$$C = \begin{bmatrix} c_r \\ c_g \\ c_b \end{bmatrix}, \quad H = \begin{bmatrix} h_x \\ h_y \\ h_z \end{bmatrix}, \quad P = \begin{bmatrix} p_r \\ p_g \\ p_b \end{bmatrix} \tag{4.5}$$

where the subscripts r, g, b denote the three channels of an RGB colour model, and x, y, and z denote the three channels of an XYZ colour model. Let $V_{c,p}$ and $V_{h,p}$ be of size (3×3) that denote the projector colour mixing matrices for the camera and human visual system, respectively. By substituting in (4.3) these matrices can be expressed as follows:

$$V_{c,p} = \begin{bmatrix} v_{rr} & v_{rg} & v_{rb} \\ v_{gr} & v_{gg} & v_{gb} \\ v_{br} & v_{bg} & v_{bb} \end{bmatrix}, V_{h,p} = \begin{bmatrix} v_{xr} & v_{xg} & v_{xb} \\ v_{yr} & v_{yg} & v_{yb} \\ v_{zr} & v_{zg} & v_{zb} \end{bmatrix} \tag{4.6}$$

From (4.2), the response of a camera to a projector, as well that of a human visual system to a projector, can be represented as

$$C = V_{c,p}P \tag{4.7}$$

$$H = V_{h,p}P \tag{4.8}$$

Existing state-of-the-art radiometric compensation methods, such as [23, 6, 31], correct for colours as seen by the calibration camera by modelling the inverse of (4.7), and do not compensate for the colours seen by a human viewer. Although these compensation techniques provide reasonably precise correction of screen uniformity, they cannot accurately calibrate to an exact desired colour. From Figure 4.1 (d), human responses to the projector primaries and white are different from the target colours as seen when calibrated

to a camera's response. This difference in colour calibration result is due to the differences in the shape of the human and camera sensor response curves used in (4.3); as a result

$$V_{c,p} \neq V_{h,p}. \tag{4.9}$$

Thus, it is evident that a mapping between the human and camera responses is required. In the following section, a mapping from camera spectral response to a model of a human spectral response using the stimulus of a projector is introduced.

4.2 Mapping of Spectral Responses

To map a given camera colour response to a human colour response, the sensitivity of both sensors to a given projector stimulus are first determined by (4.7) and (4.8). This portion of the calibration needs to be performed only once for a given projector-camera pair. The projector spectrum is used as an intermediate reference point to establish the relationship between camera and human colour responses. This relationship will then be used to eliminate any system dependency on camera pre-calibration or *a priori* knowledge of the quantum efficiency of the camera.

The mapping between camera and human responses can be determined by substituting P from (4.7) into (4.8) as follows:

$$H = V_{h,p}V_{c,p}^{-1}C. \tag{4.10}$$

To perform radiometric compensation in the colour space of content images, the human colour response H can be mapped to a colour S of size (3×1) in a standard colourspace (such as an RGB space with linear gamma) by using a known mapping, $V_{s,h}$ [28]:

$$S = V_{s,h}H \tag{4.11}$$

From (4.10) and (4.11), a mapping of the camera response into a standard colourspace can now be determined as

$$\boxed{S = V_{s,h}V_{h,p}V_{c,p}^{-1}C} \tag{4.12}$$

By transforming the colours as seen by the camera C to a camera invariant colourspace, the radiometric compensation can be performed using S, which will be referred to as the *corrected camera response*.

4.3 Compensation in Mapped Space

In this section, the relationship in (4.12) will be utilized in order to perform radiometric compensation with a camera response corrected with respect to a human viewer. This radiometric compensation will be established first for linear projector systems without ambient light. Then, to consider practical projector situations, the proposed method will be extended to systems demonstrating nonlinear projector responses and operating in the context of ambient light contributions.

Background Modeling: From (4.7), the projector intensity P for a linear system can be represented as

$$P = V_{c,p}^{-1} C. \tag{4.13}$$

In order to map the camera response to the standard colourspace, both sides of (4.13) are multiplied by $V_{s,h}V_{h,p}$:

$$V_{s,h}V_{h,p}P = V_{s,h}V_{h,p}V_{c,p}^{-1}C. \tag{4.14}$$

Due to the effect of the background, the right-hand side of (4.14) can be approximated by S (4.12), so that the required projector intensity P that compensates a background can be simplified as

$$V_{s,h}V_{h,p}P \approx S \tag{4.15}$$

$$P \approx V_{h,p}^{-1}V_{s,h}^{-1}S \tag{4.16}$$

$$P \approx KS \tag{4.17}$$

where $K = V_{h,p}^{-1}V_{s,h}^{-1}$ is of size (3×3) and represents the matrix which compensates P given a projection background and corrected camera response S.

In the case where the projector is not the sole contribution to a sensor response, as in [21], the additional light contribution can be modelled by modifying K and S as follows:

$$P = \left[\begin{array}{c|c} & k_{ra} \\ K & k_{ga} \\ & k_{ba} \end{array}\right] \left[\begin{array}{c} S \\ -- \\ 1 \end{array}\right] \tag{4.18}$$

$$\equiv K_a S_a \tag{4.19}$$

Similar to [25, 6], the expression in (4.19) can be extended to the case of a nonlinear relationship between a projector and sensor response as

$$P = F^{-1}(K_a S_a) \tag{4.20}$$

where $F^{-1}(\cdot)$ is a function representing the inverse nonlinearity of the system that can be determined empirically by exhaustively sampling the domain of the projector gamut, and then stored as a lookup table for much faster computation.

In order to model the background, given P and the corresponding S_a for a series of test colours, the matrix K_a can be obtained as the least-squares solution of (4.19) or (4.20) for the linear or nonlinear case, respectively.

Background Compensation: In the compensation phase, since the pixel colours of a given input image are provided as the target colours C to be seen on the screen and they are in the standard colourspace of S, these values are used as S to construct S_a. Next, the solved matrices K_a and S_a are used to directly compute the required projector output P for the linear or nonlinear case by using the relationship in (4.19) or (4.20), respectively.

The proposed compensation solution incorporates methods to compensate both linear and nonlinear projector-camera system, and do so by incorporating the human visual system response. This system can thus compensate any projector, regardless of the type of light source, to the correct desired colours as seen by a human viewer. Every part of these methods can be computed prior to introducing content to the projector, and has been designed for the fastest possible runtime performance. The runtime performance will be explored in section 6.0.2. This system will produce colours which are correct to a human viewer at a pixel level in the video display, provided the inputs to the system are correct, reliable and noise-free. It is therefore extremely important that the samples for each pixel model are correct, and outliers are detected and corrected.

Chapter 5

Outlier Detection and Correction

5.1 Introduction

Radiometric compensation of a projector is the process of calibrating the output of the projector to produce a colour-correct target image on a non-uniform surface. Radiometric compensation aims to achieve a more accurate representation of the projected image by modelling the non-uniformity of the system. This non-uniformity often comes in the form of a textured or coloured screen, or luminance non-uniformity of a projector source. In the context of projector-camera systems, the camera provides the necessary feedback to model the system. The industry of projection displays often requires projection on adverse backgrounds, such as the sides of buildings, walls, abstract surfaces or simply poor cinema screens [20, 3, 16, 29]. When the surfaces unavoidably have variations in colour, reflectance and texture, the imagery projected onto a given surface becomes augmented by this texture. The feedback of cameras allows radiometric compensation to be performed for most pixels with some reasonable exposure constraints, but often this modelling can fail on surfaces of extreme reflectance and produce a poor pixel model [21]. For the purpose of this study, these areas of extreme reflectance are pixels in a projector-camera system for which the calculated radiometric compensation model is unable to correctly produce visually accurate compensation of the background texture. More specifically, these are pixels which appear poorly compensated while other pixels in an image are correctly compensated.

Typically, these difficult pixels are a product of having highly specular or non-reflective surfaces, or from having noise in the form of camera sensor errors or pixel correspondence errors. Interestingly, pixel correspondence errors between the projector and camera are typically a product of the other three problems, so will not be addressed directly. It is a

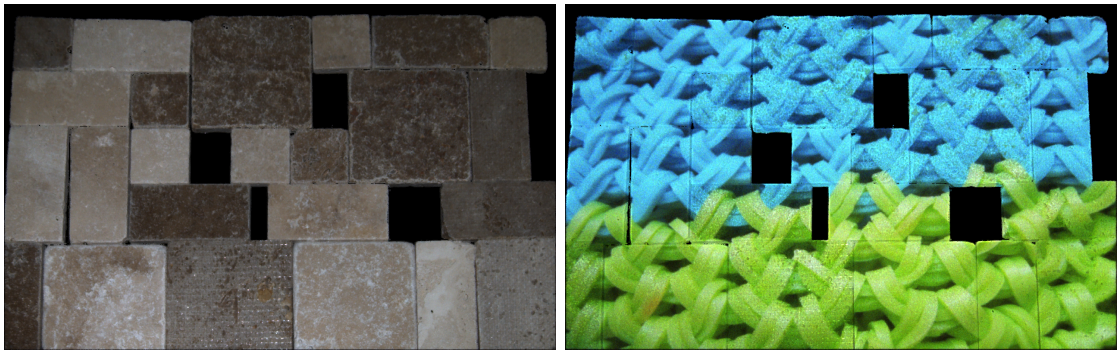

(a) (b)

Figure 5.1: (a) Sample *Stone Tile* test background and (b) compensated image using the proposed scheme. This experiment is conducted in high ambient light conditions.

challenge to not only identify these tricky pixels reliably, through a variety of projection scenarios, but also a challenge to handle these pixels with seemingly missing colour information. The aim of this study is to provide a means of rapidly identifying pixels which provide poor compensation models prior to application of the models. This would be in the form of an image mask which indicates the condition and deficiency of the compensation matrices, and by extension a mask of pixels to be handled differently than conventional system inversion solutions. Handling the poorly defined matrices can be performed either by altering the final image outputs, or by altering the values of the matrix directly. Altering the final outputs or the the value of poorly defined compensation matrices inherently becomes a denoising or inpainting problem, which lends itself well to various statistical methods. Furthermore, the mathematical derivations of this project can be generalized to systems beyond those of camera-projectors; this methodology can be applied to any multi-channel active sensor problem. All of the principles of system modelling, stability, will be derived in general for multi-channel linear systems. Exploring these principles on other systems would lie outside of the project scope.

5.1.1 Related Work

Early methods for radiometric compensation were proposed using per-pixel linear solutions of increasing complexity [2, 31, 23, 30]. For instance, in [30] a method for compensating single-channel systems is introduced, which solves for a single gain value to be applied per-pixel in each target image, thereby compensating for the luminance but not compensating multiple colour channels simultaneously. Others proposed three-channel colour compen-

sation [23, 31] by including colour mixing between projector and camera. These methods aim to produce colours to appear correct to a camera by modeling the light contributions from each projector channel to each camera channel by use of a 3×3 matrix solution per pixel. Ambient light was incorporated into radiometric compensation methods through the introduction of an additional column to the compensation matrices [31]. Recently, several methods have been proposed for radiometric compensation of nonlinear systems [16, 6, 25]. For example, a method for making the solution for per-pixel nonlinear systems feasible has been introduced in [6], where a lookup table is used to provide the compensation for each pixel. These methods still suffer on pixels where there is a structural lack of data, such as where pixels always appear black on scene, or dead camera pixels. These nonlinear methods are therefore susceptible still to tricky pixels.

These methods all create a model for the colour response of a pixel as seen by the camera when illuminated by a projector with a given RGB input. Radiometric compensation uses the inverse model to determine a required projector RGB input for a desired colour on the illuminated surface. The methods assume a well-conditioned solution exists for all pixels, which is typically not the case for regions of an image which are highly reflective or not reflective at all. In these cases, under any illumination the camera tends to see white for specular surfaces and black for non-reflective surfaces, causing a model for the pixel colours to be inaccurate. Furthermore, the values in these models vary significantly from the values in valid pixel models enough to prevent meaningful statistical operations on the entire set of solutions.

5.2 System Description

The system which is modelled and solve for radiometric compensation is the inverse projector-camera colour model for a pixel. This is the model which determines the required projector input for a desired camera output. The pixel correspondences between projector and camera are first produced using decoded structured light, such as gray-coded binary patterns. Once the geometric mapping is ascertained, the radiometric models can then be described on a per-pixel basis.

The model of a projector-screen-camera system is composed of three main transfer functions; that of each of the components. Modern projectors and cameras typically represent colour through the use of three-channel RGB values for a two-dimensional grid of pixels. The projector response $e(\lambda)$ to a given input can be modelled as a linear gain P applied to an intensity distribution as a function of wavelength $w(\lambda)$:

$$e(\lambda) = \sum_{i=1}^{c} P[i] \cdot w[i](\lambda) \tag{5.1}$$

Similarly, the camera produces a three channel response to a given light distribution:

$$C[i] = \int_{\lambda=300}^{800} r[i](\lambda) \cdot e(\lambda) d\lambda \tag{5.2}$$

When modulated by the spectral gain of a given surface $s(\lambda$, the following relationship between a projector and camera is developed:

$$C[i] = \int_{\lambda=300}^{800} s(\lambda) \cdot r[i](\lambda) \cdot \sum_{j=1}^{c} P[j] \cdot w[j](\lambda) d\lambda \tag{5.3}$$

Rearranged we see $C[i]$ is thus a sum of products:

$$C[i] = \sum_{j=1}^{c} P[j] \cdot \int_{\lambda=300}^{800} s(\lambda) \cdot r[i](\lambda) \cdot w[j](\lambda) d\lambda \tag{5.4}$$

This can be represented as:

$$C[i] = \sum_{j=1}^{c} P[j] \cdot v_{[i,j]} \tag{5.5}$$

$$V = \begin{bmatrix} v_{RR} & v_{RG} & v_{RB} \\ v_{GR} & v_{GG} & v_{GB} \\ v_{BR} & v_{BG} & v_{BB} \end{bmatrix} \tag{5.6}$$

Each component is linear, giving the 3 system V, plus an ambient light component α:

$$C = VP + \alpha \tag{5.7}$$

The system inverse compensation is thus:

$$P = V^{-1}C + A \tag{5.8}$$

where V^{-1} is the inverse system and A is the ambient offset. The inverse system is modelled rather than known, and the intent of this study is to analyze and improve the stability of this system.

5.2.1 Stability Analysis

The system inverse model is created by solving for the least-squares solution to an over-constrained system V using a series of projected test colours. For each pixel, a compensation matrix V^{-1} must be solved using the projector and camera values P and C for n test colours:

$$P \simeq \hat{V}^{-1}C \tag{5.9}$$

where P is an $n \times 3$ matrix, and C is an $n \times 3$ matrix.

In many cases, the solution provided may not be meaningful or useful. This typically occurs for black regions or highly specular regions, where a change of projector intensity is not reflected in the camera response. In such cases, the solution for the system inverse $\hat{V}^{-1}$ is rank-deficient, and results in many 0-valued elements to compensate for specular pixels, or *inf*-valued elements to compensate for black (non-reflective pixels. This can be seen in Figure 5.1 where there are both specular and black regions in the stonework background. Such poor conditioning of the forward and inverse system produces colour values which lie significantly outside the range of the rest of the compensated image values. This results in such values being discarded or clipped to fit within the range of possible projector outputs. Additionally, these outlier values impede the use of any global normalization of the projector outputs. Identifying these difficult pixels would improve the use of normalization and global techniques, and result in the improvement of overall projecting image quality.

Stability and conditioning analysis can be used to identify pixels which respond unfavourably to compensation, and potentially improve compensation results. If the pixels are properly identified prior to compensation or calculation of the inverse system, the pixels can be handled with a different compensation scheme.

5.3 Proposed Method

The proposed method for correcting poor models for pixelwise radiometric compensation can be decoupled into two major steps. Firstly, the model accuracy of the pixels must be quantified, and from this measurement the poor pixel models must be identified. The method classifies pixels into two categories: those which provide reliable enough models to provide visually correct radiometric compensation, and those which cannot. Certain pixels might provide difficult projection backgrounds, such as dark spots and blemishes

on a screen, yet still provide well-posed compensation solutions. Here, the aim is to identify pixels which do not provide well-posed matrix solutions for their compensation, and assuming some continuity of projection surface with neighboring pixels. pixels which do not provide a well-posed matrix solution for compensation may be a result of dead camera pixels, screen defects smaller than a projector or camera pixel, or poor correspondences (mapping) from camera pixel to projector pixels due to errors in the structured light. The proposed method identifies poor pixel models by computing the error of the system models with respect to the measurements used to generate the models. Pixel models which poorly predict their own measurements are identified by calculating the mean squared error between the actual and predicted projector images used to compute the pixel models. Secondly, following the identification of poor pixel models, the models must be filled in from the models and measurements of the surrounding pixels.

5.3.1 Stability Analysis and Outlier Detection

In the proposed method, the accuracy of the radiometric model for each pixel is quantified to provide means of identifying poor pixels. The model for a particular pixel is generated from a series of projector inputs P and corresponding camera outputs for the same pixel on surface, C. When the projection surface is very dark or specular, the response of the camera does not respond correctly to the light from the projector. In these cases, the model is unstable and is potentially rank-deficient, which The accuracy of a particular inverse pixel model is measured by its ability to correctly predict the required input to a projector, so the accuracy of the models is calculated from the mean squared error of the inverse models when used to generate the projector input from the training set. This leads to a simple and reliable way to diagnose poor models and characterize the accuracy of well-posed inverses, by calculating the prediction error on the training set:

$$\epsilon = \frac{\sum_n^{i=1}(\hat{P}_i - P_i)^T(\hat{P}_i - P_i)}{n} \tag{5.10}$$

$$\epsilon = \frac{\sum_n^{i=1}(KC_i - P_i)^T(KC_i - P_i)}{n} \tag{5.11}$$

This mean squared error can then be used to directly determine poor pixels by setting a threshold based on the maximum acceptable perceived error:

$$\epsilon > \hat{\epsilon} \tag{5.12}$$

5.3.2 Infilling

After the identification of poorly modelled pixels, the model for the pixel must be filled in using the remaining information nearby. This can be done either through filling in poor pixels in the training image set, then solving for the pixel model again, or by directly interpolating the systems. The radiometric compensation model for a particular pixel in a projector image, corresponding to a given camera image can be represented as:

$$P_{D([x,y])} = K_{[x,y]} C_{[x,y]} \tag{5.13}$$

Where $[x, y]$ are the pixel coordinates, and $D()$ provides the correspondence mapping between projector and camera pixel space. The pixel model K is typically constructed from training sets consisting of corresponding P and C values, but in the case of poorly modelled pixels, is constructed from information from surrounding pixels. This leads to two separate paths for interpolating the missing pixel models, either by interpolating the systems directly, or by interpolating the training set used for generating the models. In the case of radiometric compensation, square 3 by 3 matrices are used for the pixel models, each of which is determined here using the least-square-error moore-penrose pseudoinverse for over-constrained systems. The use of matrix inverses of any size leads to the property that interpolating the inverse systems is not equivalent to inverting an interpolated set, which is evident from even 2 by 2 matrix systems:

$$(\alpha A + \beta B)^{-1} = \begin{bmatrix} \alpha a_1 + \beta b_1 & \alpha a_2 + \beta b_2 \\ \alpha a_3 + \beta b_3 & \alpha a_4 + \beta b_4 \end{bmatrix}^{-1} \tag{5.14}$$

Which is nonlinear in both A and B, as shown by [22] where the general inverse of sums is derived:

$$(\alpha A)^{-1} - \frac{1}{1 + trace((\beta B)(\alpha A)^{-1})} (\alpha A)^{-1} (\beta B)(\alpha A)^{-1} \tag{5.15}$$

The proper inverse solution of a linear combination of matrices can therefore never be modeled by a linear combination of their respective inverses, meaning that the matrix resulting from interpolating the neighboring inverse matrices K is not the same as the result by interpolating C and P, then solving a new inverse model. The consequence of this is that interpolating the pixel models is only valid if using nearest neighbor interpolation, and so any interpolation used to determine the pixel models must be performed on the input images from which they are calculated. The pixel model is then be computed from the interpolated input images using a thin-plate spline interpolation.

The interpolation of the input (camera) images is performed on the pixels which have been previously identified as poorly modeled. In the sequence of input images from the camera, each poorly modelled pixel is interpolated from the neighboring well-modelled pixels surrounding it. This is performed independently for each camera image, but the projector output does not need to be interpolated as it is known for all pixels.

The interpolation method selected was the thin plane spline, to provide second-order smoothness constraints as found in [4]. The thin-plate spline constraints kernel is as follows:

$$Q = \begin{bmatrix} 0 & 0 & 1 & 0 & 0 \\ 0 & 2 & -8 & 2 & 0 \\ 1 & -8 & 20 & -8 & 1 \\ 0 & 2 & -8 & 2 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix} \tag{5.16}$$

The pixels with poor models are interpolated using a least-squared solution enforcing a TPS constraint for a given image, where the outlier pixels are discarded from the calculation. The interpolated values of the pixels are the only ones which need to be solved; the remaining good pixels are retained from the original image and are not smoothed during interpolation. There is the potential here however to recalculate the value of all pixels in a given input image based on smoothness constraints and weights based on the pixel model quality, though the only real advantage of this is general noise reduction and image smoothing in camera images, which risks softening crucial edges.

5.4 Quantitative Performance

The proposed method was quantitatively assessed on both real and artificially constructed compensation scenarios. The tasks of both identifying poor pixel models and correcting for the models are evaluated separately, using the same dataset. The accuracy of infilling the pixel models could only be assessed by using an image with ground truth, thus requiring the use of an artificial dataset. In these experiments, the training set consists of the flat-field colour images used as projector inputs,and the corresponding camera images of the projection surface under this illumination. These images were all taken of backgrounds without dark or specular pixels, and naturally provide good colour models for each pixel. To create a set of images with poor pixels, a series of real-world radiometric compensation training sets were augmented with noise. Noise was added solely to the camera images of a training set, with a specific proportion of randomly selected pixels having added noise. To simulate regions of steep reflectance in the images, the pixels selected for added noise had

either salt or pepper noise consistently applied across all camera images. The peppered pixels emulate pixels of unusually dark response in the camera, whereas the salt noise mimics specular regions. The mask of noisy pixels was retained for the noised dataset, to assess the accuracy and sensitivity of the method in identifying poor pixels. The proposed method was compared with the ground truth and other methods were not implemented.

A dataset comprised of 4 series of 12 training images each is used to test the ability of the proposed method to detect pixels with poor models. In each series of 12 images, the images are all identically aligned on the projection surface, and so all share identical pixel correspondences with the projector image. For each series, 2% of pixels were randomly selected to have noise applied to them, and these pixels then had their RGB values replaced with random values between 0 and 255. The same pixels were perturbed in all images of a series, but each series had a different random selection of pixels. To reduce the statistical impact of any particular pixels, the entire process of selecting, detecting bad pixels and replacing them was repeated 100 times per image series.

5.4.1 Detection of Outliers

The outlier detection scheme was evaluated quantitatively on the artificially noisy datasets, and verified qualitatively on real images of compensation scenarios. The dataset of 4 series of 12 images each was augmented using salt and pepper noise, and the mask of noisy pixels provides the ground truth for outlier detection. On each dataset, the pixel models were determined, then the accuracy of these models were assessed using the proposed outlier detection scheme. This process can be seen in Figure 5.2. This figure shows one of the original images from the Rainbow background, the noise mask used to perturb the base image, the colour noise that is used to replace pixel values, and the final noisy image.

The mask of pixels found to be outliers or noisy was then compared to the ground truth. The accuracy was measured by calculating the percentage of correctly identified inlier pixels (ones without noise added), and the percentage of correctly identified outlier pixels (pixels with noise added). This can be found in Table 5.1. The method performed better on the smooth, consistent backgrounds such as the rainbow background and parchment background, whereas the less stationary backgrounds of stone and brick posed a greater challenge. Overall, the method correctly classified 95.6% of inliers, and correctly classified 98.4% of outliers in the worst scenario, the stonework surface.

This surface was chosen as an example of the pixel mask which is generated by the proposed outlier detection scheme. In Figure 5.3, the mask of good pixels can be seen in white, where all pixels which do not provide a strong model are black. As can be seen in the

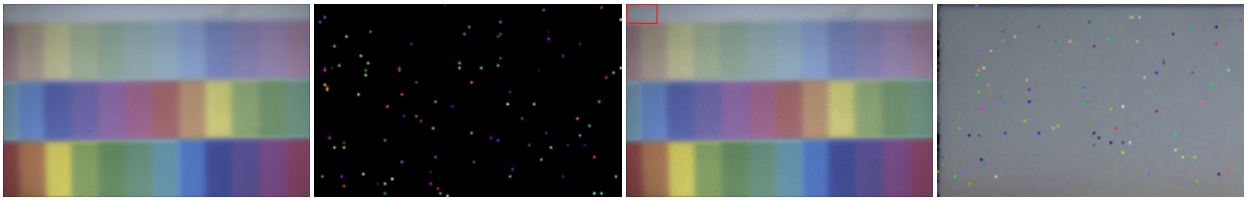

(a) (b) (c) (d)

Figure 5.2: An illustration of the noise applied to a captured dataset with original image (a), the mask of colour noise to be applied to the image (10x zoom, ROI in next image) (b), an image with noise applied values with which the pixels will be replaced with zoom ROI, and finally (d) shows the result of noise applied to the image, zoomed for detail.

figure, all pixels which are not illuminated by the projector are considered outliers. This is because the pixel values which generated the models are completely unable to predict the lightsource intensity illuminating them. Visually speaking, the pixel mask clearly includes all illuminated pixels and excludes all non-illuminated pixels, which also includes the portions of the illuminated background which are simply too dark to reflect sufficient light back to the camera. It can also be seen that certain highly specular pixels have been discarded, as they similarly provide no ability to predict their illumination colour, but for a different reason.

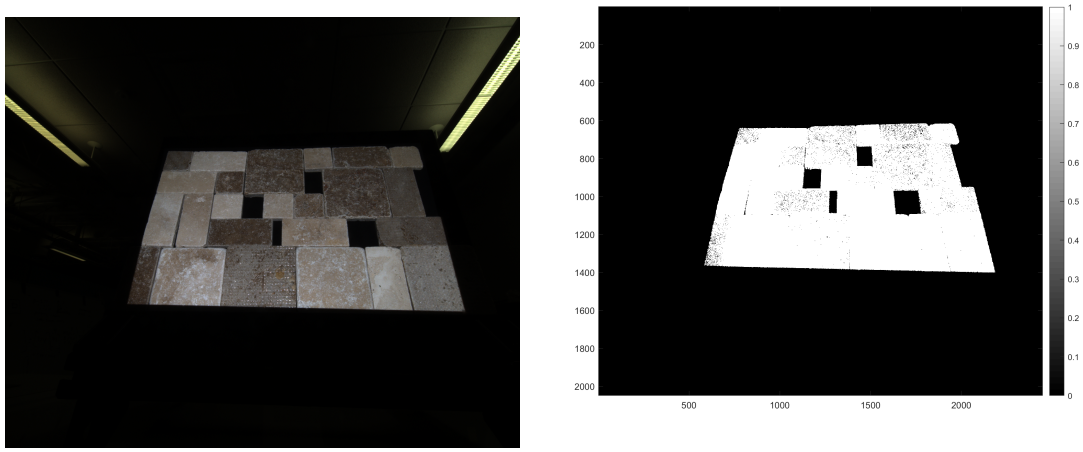

(a) (b)

Figure 5.3: Original Stonework Image Capture under White Light (a), and Pixel Mask (b)

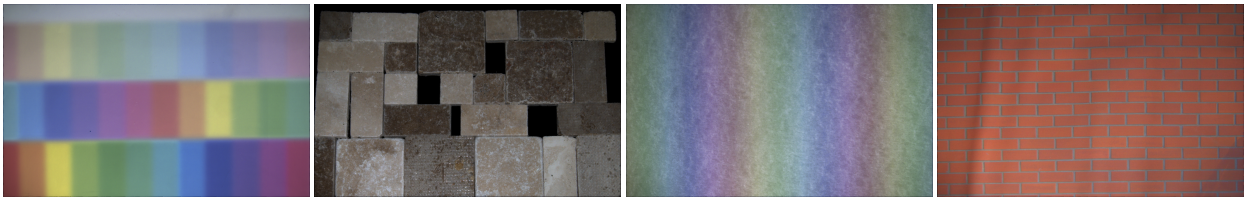

(a) (b) (c) (d)

Figure 5.4: An illustration of the backgrounds used in the dataset with Rainbow (a), Brick (b), Parchment (c), and Brick (d). All have been warped to projector pixel coordinates.

5.4.2 Accuracy of Correction

The accuracy of the correction was assessed in two parts, namely the pixel error and the PSNR. Both of these error metrics were measured between the original non-noisy camera images and the images with noisy pixels filled in with interpolated values. The noisy and corrected images can both be seen in Figure 5.5, where the rainbow background allows the speckle of outliers to be seen more clearly. Visually, this infilling in nearly indistinguishable from the original image, as the human eye is very forgiving to slight single-pixel differences. The performance of the infilling method is assessed in Table 5.1, where the mean pixel error is anchored in perspective by the prediction error of the pixel models of unperturbed pixels. The PSNR overall correlates with the mean pixel error, simply by the nature of the PSNR calculation. The PSNR is generally quite high on all backgrounds as only 5% of pixels were perturbed with noise. The accuracy of the correction is demonstrated by these error metrics, and is relatively low on the smooth backgrounds of the rainbow and parchment textures. Comparatively, the error is much higher on the non-stationary backgrounds of the stone and brick, where the texture has much steeper gradients and neighboring pixels have less relationship to the pixel being infilled.

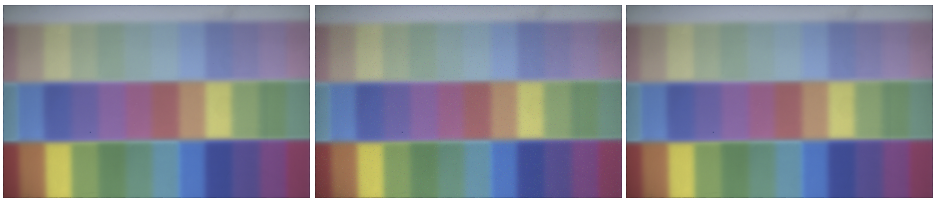

(a) (b) (c)

Figure 5.5: Original Rainbow Image (a), noisy image (b), and denoised image (c)

Image Series	Mean Error	Mean Error (Noisy Pixels)	PSNR	Correctly Classified Inliers	Correctly Classified Outliers
Rainbow	0.04	12.80	57.01	99.2%	98.8%
Stone	0.13	40.69	48.99	95.6%	98.9%
Parchment	0.06	18.05	55.51	98.7%	98.4%
Brick	0.08	24.36	55.76	98.5%	98.6%

Table 5.1: Comparison of outlier Detection and Infilling accuracy among datasets

5.4.3 Discussion

Through the quantitative results, is has been shown that projector pixels which provide poor compensation models and consequently poor compensation results, can be accurately identified and inferred from neighboring strong pixels. The method for classifying poor pixels demonstrated a larger variance in accuracy between correctly classifying noise-free pixels when compared to the noisy pixels. The percentage of correctly classified inliers varied from 95% to nearly 99%, but the percentage of correctly classified outliers varied from 98.4% to 98.9%. This indicates there are separate mechanisms for the wrong classification in each case. When investigating the wrongly classified inliers (inliers classified as outliers), these so-called inliers were in fact difficult pixels in the original camera images. These images were taken from real camera captures of a live background scene, which includes dark pixels in the case of the stonework, and natural camera sensor noise in all portions of the dataset. The outlier detection technique was finding outliers inherent in the dataset that were not added in the noise process, as the ground truth of outliers did not account for existing native outliers in the dataset.

More interesting is the remarkably consistent classification rate of the outliers which were introduced in the noise process. This hints at a structural problem in the identification process, but is in the method is which the noise was generated. The noise that is used to perturb the images is RGB random noise in specific pixels, which have a small chance of randomly forming a pixel solution which is consistent. In other words, even a 14 by 3 matrix of random values can form a consistent pattern a small percentage of the time. As these randomly generated noisy pixel values are simply fed into a matrix solver and the prediction error is thresholded, these noisy pixels can slip by the detector in cases where they happen to be consistent and predictable from the system model matrix they produce.

5.5 Conclusions

The proposed method offers reliable, efficient means to both identify pixels in a projector-camera system which have been modeled poorly, and means to infill the systems for these poor pixels by using the information from surrounding pixels. The method was demonstrated to successfully detect over 98% of outliers in a dataset of images in a radiometric compensation run. The method was also shown to discriminate outliers from reliable pixels very well, with over 95% of reliable pixels classified correctly of even significantly adverse backgrounds. The method of interpolating data prior to recomputing pixel models has also been shown to be very effective, achieving a mean pixel RGB error of 12.8 on relatively smooth backgrounds, and a mean pixel RGB error of 40.687 on adverse backgrounds while still appearing visually consistent.

Further study is necessary into enhancing the interpolation techniques used for inferring the missing pixels in images, particularly in images or other data where textural patterns can emerge. Such enhanced interpolation techniques to be tested may include generative neural nets or classical textural inpainting techniques. Further study is also warranted into the noise processes used for creating artificial datasets, as to provide true outliers which do not provide good models for a system solver.

Chapter 6

Results

The results that follow are in three parts: demonstration of compensation of nonlinear projectors, compensation using the Human Visual Space, and outlier removal.

For projector calibration, eight flat-field test patterns are used: white, black, three primaries (red, green, blue), and three secondaries (cyan, yellow, magenta). The use of additional colours, beyond the commonly-used three primary and black colours, leads to greater stability of the solution for each pixel. The camera captures of each of these flat fields and the projected colours are then remapped using the matrix from (4.20) to determine the least-squares solution to K_a, where an RGB colourspace with linear gamma was chosen for the camera output.

6.0.1 Qualitative Results

Several backgrounds and target images have been selected to test the capabilities of the proposed radiometric compensation method under several challenging conditions. In particular, the proposed radiometric compensation scheme is evaluated by using three challenging background surfaces, namely *Brick, Rainbow, Textured Rainbow*, and the compensation is applied to nine different target images, namely, *Astronaut, Car, Cubes, Flower, Balloons, Skating, Plate, Yellowstone* and *Waterfall*. The first and second columns of Figure 6.1 show the test backgrounds and the target images, respectively. The *Brick* background represents the scenario of projection onto brick buildings, which is very challenging, particularly when compensating for the dark borders around bricks. The *Rainbow* and *Textured Rainbow* backgrounds both include a full range of colours, with *Rainbow* designed to have a

range of extreme but smooth colour patches, whereas *Textured Rainbow* also requires compensation for higher frequency texture patterns in addition to the variations in chroma.

Figure 6.1 also demonstrates sample projector outputs and compensated results for the proposed scheme, and the captured results for projection without radiometric compensation and projection with radiometric compensation but without HVS. The first row of Figure 6.1 shows the results of projecting the *Astronaut* target image on *Brick* background. It is clear that the proposed scheme can provide compensation quality better than that of projection without using HVS, as the proposed scheme compensates to the correct colours and has reduced the effect of edges in the final image. In the second and third rows, the *Car* and *Cubes* target images have been projected; the proposed scheme offers colour uniformity, non-noticeable background edges, and maintains visual quality for the human viewer.

The middle three rows of Figure 6.1 show the results of projecting *Flower, Balloons* and *Skating* images on the *Rainbow* background. Unlike the method of projection without HVS, the proposed method provides not only highly uniform compensated images, but also shows accurate colours that are very close to the content colours. This also confirms the ability of the proposed scheme to solve for highly colourful background and content scenarios, and for the capability of projecting a sample sports content on a challenging background.

Finally, the last three rows of Figure 6.1 illustrate the results of projecting *Plate, Yellowstone* and *Waterfall* content images on the *Textured Rainbow* background. These results demonstrate the ability of the proposed method to produce compensated colours that are closer to the target colours of an image, rather than that of compensation without using HVS. This performance is consistent across the various textures, colours and environmental conditions.

6.0.2 Quantitative Results

For purposes of practical evaluation, two different projector classes[1] are used to assess the uniformity and accuracy of the compensation. The camera used in the system is a 5MP camera[2], which is used to provide feedback to the system for compensation. The projector models differ in their light sources, and include an LED based projector of linear colour response and a mercury lamp based projector of nonlinear colour response. The *Rainbow*

[1]The projectors used were Christie Matrix StIM (LED) and Christie DWX600-G (Mercury Lamp).
[2]The PointGrey Flea3 Camera with GigE was used.

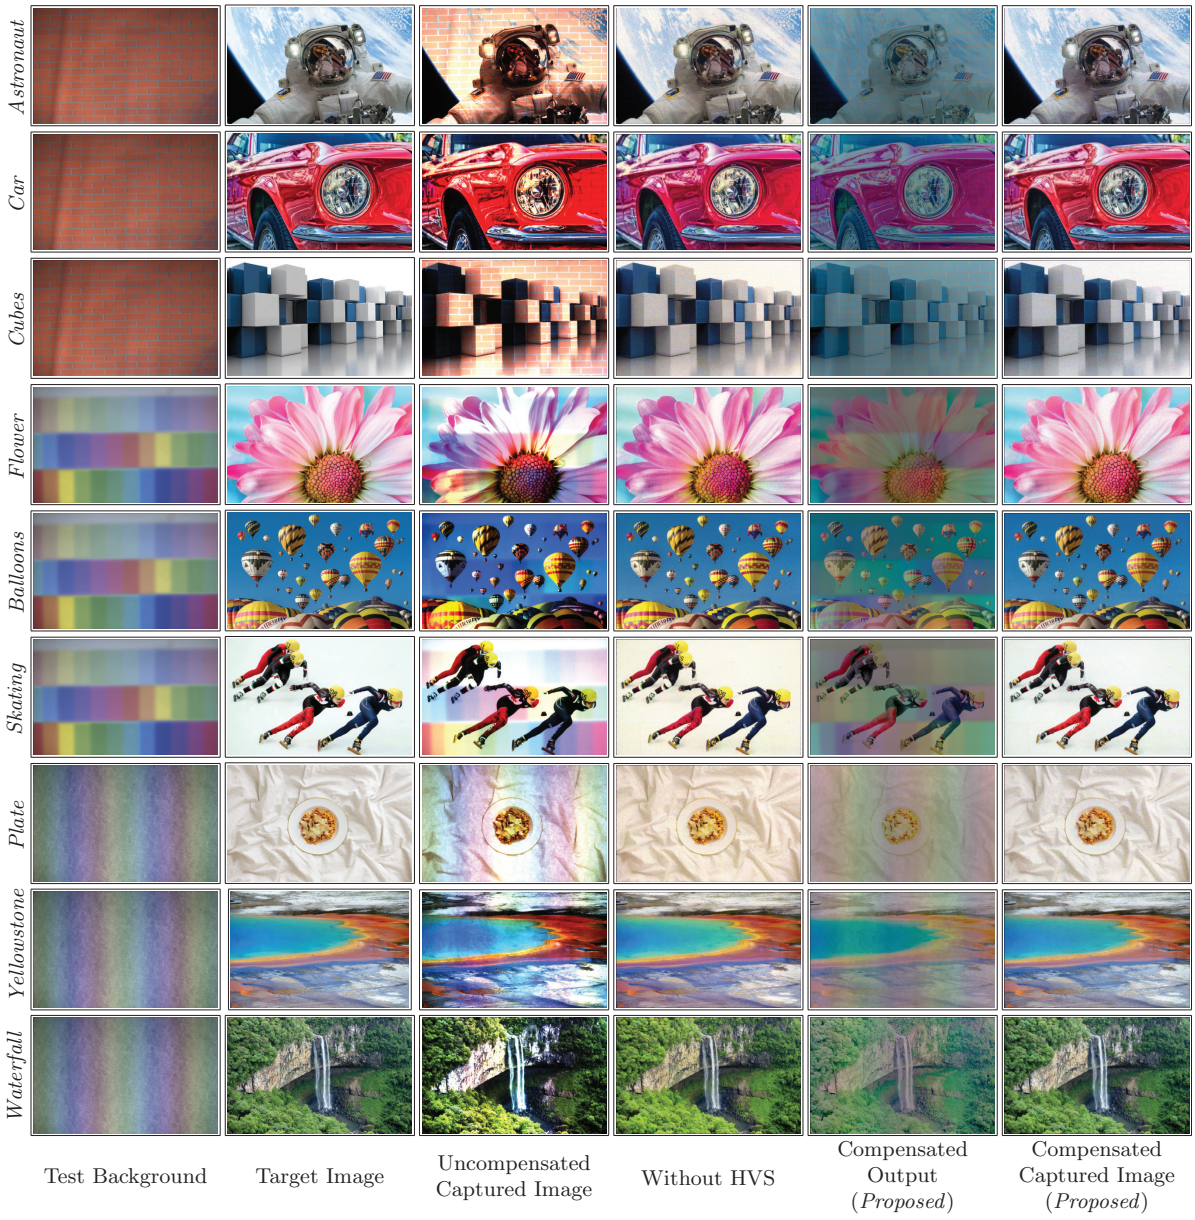

Figure 6.1: Sample qualitative results for the proposed scheme, uncompensated projection, and compensation without HVS correction on *Brick* (1st, 2nd and 3rd rows), *Rainbow* (4th, 5th and 6th rows) and *Textured Rainbow* (7th, 8th and 9th rows) backgrounds using nine different test images. All images are shown in projector coordinates.

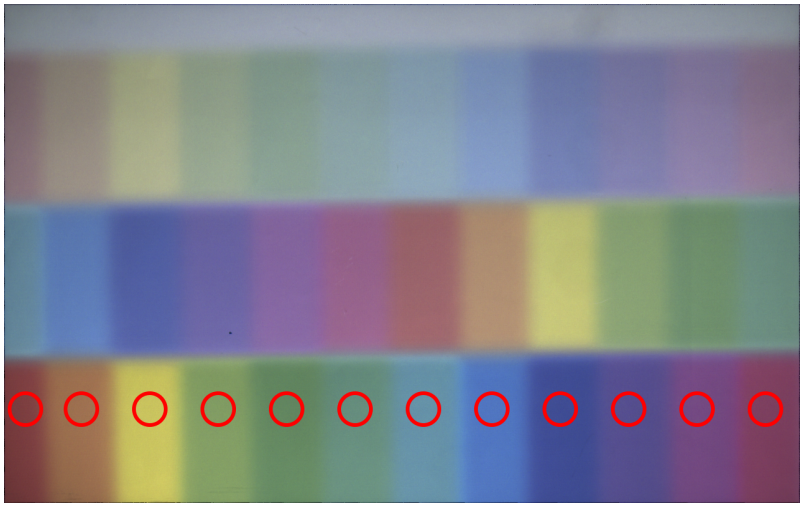

Figure 6.2: Illustration for the twelve different sample locations on the *Rainbow* background used for the spectroradiometer measurements.

background pattern shown in Figure 6.1 is used to provide a challenging background with uniformly coloured patches for which the compensation accuracy is measured by a spectroradiometer. This test background is designed with twelve different background colours, and each colour either being a primary, secondary, or tertiary colour in the RGB and CYM colour systems to challenge the limits of the compensation methods. For each compensation method, a spectroradiometer[3] outside of the calibration loop is used to measure the screen compensation for four target colours on the coloured background patches. Each background patch provides a different uniform region upon which to position the circular measurement area of the spectroradiometer as seen in Figure 6.2, where the spectroradiometer measures the incident light from a 5 degree cone.

Measurements of xy and Luminance: From Figure 6.3, it can be seen that the proposed method significantly improves the white point in the resulting compensated images. The proposed method accurately compensates different projectors to the same reference white as seen by a human viewer, resulting in colours on screen which are perceptually consistent with a given target image. In addition, by compensating the projector to the correct white, the resulting compensation is more uniform overall for white, indicating not only greater fidelity for white images is achieved, but also the final image on screen is more uniform overall. Figure 6.3 also shows a fundamental limit of this method. This method can be used to accurately set the white point for compensation, but the chroma

[3]This test was carried out using the JETI Specbos 1211.

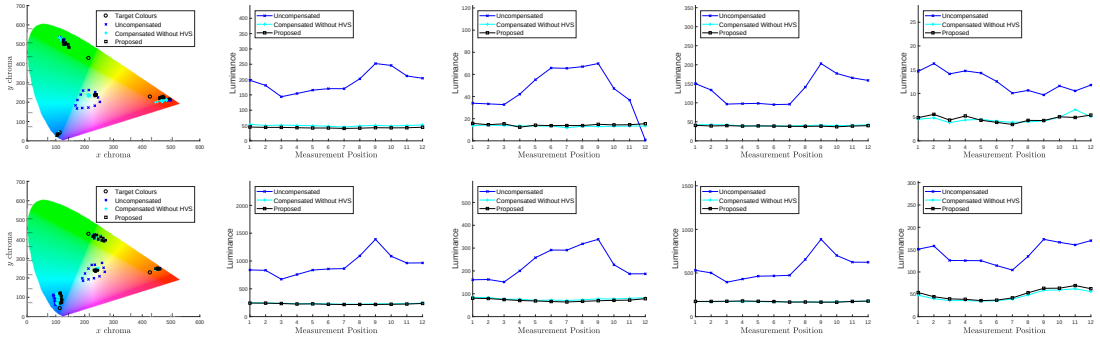

Figure 6.3: Comparing the measurements of the proposed method in the xy colourspace before and after compensation with and without using HVS, and the corresponding luminance plotted for each target colour, where a Christie Matrix StIM (LED) and DWX600-G (Mercury Lamp) projectors are used in the 1^{st} and 2^{nd} rows, respectively. In this figure, greater uniformity is indicated by a tighter points cluster in the xy colour chart or a flatter line in the luminance plots.

of the compensated primary colours remains anchored near the original values. This method performs accurate compensation of a white field on strongly coloured background, but tended to keep the primary colours of red, green and blue anchored near their uncompensated positions while compensating using luminance alone. This provides the primary avenue of further work; an investigation into a method for accurately producing all desired target colours on a projection surface with a variety of light sources. Despite this, as shown in Figure 6.1 the proposed method is shown qualitatively to compensate for extremely adverse surfaces with minimal edge artifacts incurred from the backgrounds.

The last four columns of Figure 6.3 show the luminance required to create flat colour fields of white, red, green and blue that appear uniform using this method. As seen in this figure, the luminance drop is comparable to compensation without HVS mapping, and the flatness of these curves demonstrates the luminance uniformity of the image. This difference in luminance is exactly proportional to the amount of light the most troublesome regions of the image are capable of reflecting, as the entire compensated image must be reduced to have the same colour and luminance of the least reflective (darkest) portions of a background. Similar to radiometric compensation without human colour mapping, the proposed method cannot increase the light output capabilities of a projector, but instead it must reduce the amount of light on screen to produce perceptually uniform images as seen in Figure 6.3.

Compensation Error: In order to evaluate the error between a target colour and the compensation measurements of a given method, the CIE ΔE [28] is used. This metric is used to compute colour distance between a colour seen on screen and the desired target colour. The target colour provides only a chroma target for the compensation, and the luminance is dependent on the output capabilities of the projector in the system, thus the luminance component of the ΔE must be determined separately. For this investigation, the desired luminance is taken to be the average luminance of all screen patches, as deviation from this luminance indicates non-uniformity of the compensated image. In this way, the distance to both target chroma and the overall screen luminance are combined into a single error metric to quantify both accuracy and consistency of the compensated projector output.

Table 6.1 shows a comparison of the mean and standard deviation of the ΔE values obtained for the proposed method, performing radiometric compensation projection without HVS, and projection without compensation. As in the previous experiment, white, red, green and blue are used as the target colours using the same two types of projectors. As shown in this table, the proposed method offers the lowest average ΔE compared to that of radiometric compensation without HVS in six out of eight test cases, while maintaining nearly the same performance as non-HVS compensation in the remaining two cases.

Table 6.1: Comparing means and standard deviations of ΔE for uncompensated projection, compensated projection but without using HVS and the proposed method, where the ΔE values are calculated for twelve patches when targeting four different colours. The ΔE [28] is a function representing the colour distance between the spectrometer measurement of a given patch and the given target colour. Note: boldface and underscore denote the best and second best results, respectively.

Methods	Christie Matrix StIM (LED) Projector				Christie DWX600-G (Mercury Lamp) Projector			
	White	Red	Green	Blue	White	Red	Green	Blue
Uncompensated	64.68 ± 29.00	86.12 ± 40.93	149.10 ± 52.1	66.53 ± 27.45	78.24 ± 26.91	68.08 ± 23.74	68.02 ± 35.00	265.37 ± 112.51
Comp. without HVS	$\underline{16.52} \pm 5.28$	$\underline{23.32} \pm 10.43$	$\underline{80.05} \pm 26.36$	$\underline{32.68} \pm 12.60$	$\underline{13.78} \pm 5.18$	$\mathbf{30.81} \pm 9.89$	$\mathbf{40.98} \pm 23.01$	184.62 ± 71.02
Proposed Method	$\mathbf{2.40} \pm 1.38$	$\mathbf{10.56} \pm 7.64$	$\mathbf{68.28} \pm 21.91$	$\mathbf{32.32} \pm 4.97$	$\mathbf{5.42} \pm 2.97$	$\underline{31.23} \pm 10.20$	$\underline{41.56} \pm 23.22$	$\mathbf{122.98} \pm 56.10$

Computational Cost: The storage requirements and computational time of the proposed method are verified through a Matlab implementation of the proposed method[4], and then compared to that of the existing nonlinear method in [6]. The memory and computation times for nonlinear projector systems are assessed for three different common projector resolutions (720p, 1080p and 4K). As shown in Table 6.2, the proposed method uses less than $\frac{1}{5}$ of the memory and is 13× faster in compensation time than that of the method

[4]A CPU with 4 cores operating at 3.40GHz was used.

Table 6.2: Comparison of the memory and the time complexities of the proposed scheme with and without using HVS mapping, and that of the method in [6] at three different resolutions. Note: boldface denotes the best results and * denotes a calculated value.

Methods	Memory (GiB)			Time, CPU (s)		
	720p	1080p	4K	720p	1080p	4K
Grundhfer & Iwai [6]	2.50*	-	-	7.1	-	-
Comp. without HVS	**0.46**	**0.56**	**1.17**	**0.53**	**0.82**	**2.10**
Proposed Method	**0.46**	**0.56**	**1.17**	**0.53**	**0.82**	**2.10**

in [6], while maintaining the same computational complexity of radiometric compensation without HVS. The proposed scheme requires only a modest increase in time (less than 1 second per 4K image) during the image capturing portion of the background modelling phase than that of the baseline, due to correcting the captured camera colours (4.12).

Chapter 7

Conclusion

7.1 Summary of Thesis and Contributions

In this thesis, a comprehensive method for radiometrically compensating projector-camera systems for a human viewer has been presented. The method consists of two distinct parts: calibrating nonlinear projector systems, and calibrating the system to colours specified by a human visual model.

The proposed method solves the nonlinear radiometric compensation problem by separating the linearization problem from that of linear radiometric compensation. Experimental results have shown that the proposed scheme has allowed compensation to be applied with a fraction of the memory and computational requirements than that achieved by a recent state-of-the-art method. Additionally, the proposed scheme has offered a very low compensation error when evaluated by a sensor outside of the compensation feedback loop. Compensation of an adverse multi-coloured background was accomplished with a maximum error of 5.13 over the field of the image. This work will be extended to video applications, and to eliminate the dependence on camera colour response for the final compensation.

It has been shown that by properly modelling the colour response of a human viewer, this method can produce approximately the same colours on a surface using different light sources and cameras. This eliminates the need for cameras and projectors which are previously radiometrically calibrated, with camera colour mappings instead being determined from the known spectral output of a projector. Experimental results have shown that the proposed method with using HVS generally offers the lowest average radiometric compensation error and closer to the target colours than that of compensation without HVS.

The direct consequence of this method is the ability to calibrate and match camera colour responses to others, or alternatively colour match two projectors without the need to be seen by the same camera. Further work is planned to investigate the impact of this method in performing colour calibration of blended screen multi-projector displays. Extending the proposed colour correction with respect to a human viewer to other nonlinear radiometric compensation methods is also a potential extension to this work.

7.2 Future Work

7.2.1 Compensation with Multiple Viewpoints

The use of radiometric compensation is applicable to large-scale, blended projector displays on complex surfaces such as buildings and statues. These applications typically require multiple cameras in addition to the projectors to completely cover the surface. The proposed method has been shown to correctly compensate projector surfaces to a single viewpoint, using a single camera. Future work will include the use of multiple camera images with partially shared surface coverage in order to compensate a seamless multi-projector display too large to be covered by a single camera. This work will involve blending camera images to negotiate differences in observed luminance and colour between the viewpoints.

7.2.2 Establishing Generalized Projector Models

The proposed method for radiometrically compensating projector colours to the human visual space has been effectively demonstrated using measured spectral models of the projectors using a spectroradiometer. It has been observed and that certain projector classes are internally consistent in spectral outputs when different projectors are compared, leading to the potential of sharing a single spectral profile between projectors. Future work will include developing colour models for different projector classes, with different light sources, allowing the method proposed in Section 4 to be used without the need of per-projector modelling. The manufacturing uniformity of projectors, and in particular their light sources can potentially lend itself to generalized models of certain projector classes. This would allow multiple projectors within the same projector class to be compensated using the same profile, thus removing the need for each projector to be characterized at the time of manufacture, reducing both cost and time required for the process.

References

[1] M. Ashdown, T. Okabe, I. Sato, and Y. Sato. Robust content-dependent photometric projector compensation. In *Proc. IEEE CVPR Workshops*, pages 1–6, 2006.

[2] Oliver Bimber, Daisuke Iwai, Gordon Wetzstein, and Anselm Grundhoefer. The visual computing of projector-camera systems. *Computer Graphics Forum*, 27(8):2219–2245, 2008.

[3] M. Brown, A. Majumder, and R. Yang. Camera-based calibration techniques for seamless multiprojector displays. *IEEE Trans. on VCG*, 11(2):193–206, 2005.

[4] P. Fieguth. Multidimensional models: Deterministic models. In *Statistical Image Processing and Multidimensional Modeling*, pages 149–154, October 2010.

[5] A. Grundhöfer. Practical non-linear photometric projector compensation. In *Proc. IEEE CVPR Workshops*, pages 924–929, June 2013.

[6] A. Grundhöfer and D. Iwai. Robust, error-tolerant photometric projector compensation. *IEEE Trans. on Image Process.*, 24(12):5086–5099, 2015.

[7] Anselm Grundhofer and Oliver Bimber. Real-time adaptive radiometric compensation. *IEEE Trans. on VCG*, pages 97–108, 2008.

[8] T. Huang, C. Kao, and H. H. Chen. Quality enhancement of procam system by radiometric compensation. In *IEEE Int. Workshop on MMSP*, pages 192–197, 2012.

[9] T. Huang, T. Wang, and H. H. Chen. Radiometric compensation of images projected on non-white surfaces by exploiting chromatic adaptation and perceptual anchoring. *IEEE Trans. on Image Process.*, 26(1):147–159, Jan 2017.

[10] Tai-Hsiang Huang and Huanting Chen. Radiometric compensation for ubiquitous projection, 2013.

[11] Q. Jia, H. Xu, J. Song, and X. Gao. Research of color correction algorithm for multi-projector screen based on projector-camera system. In *Proc. ISDEA*, pages 1285–1288, 2012.

[12] Deane B. Judd. Extension of the standard visibility function to intervals of one millimicron by third-difference osculatory interpolation. *J. Opt. Soc. Am.*, 21(5):267–275, May 1931.

[13] Sanjeev J. Koppal. *Lambertian Reflectance*. Springer US, Boston, MA, 2014.

[14] M. Lee, H. Park, and J. Park. Fast radiometric compensation accomplished by eliminating color mixing between projector and camera. *IEEE Trans. on Consumer Electronics*, 54(3):987–991, 2008.

[15] Jen-Shuo Liu and H. H. Chen. Preserving image color appearance on non-white projection surfaces. In *IEEE Int. Conf. on Multimedia Expo Workshops*, pages 1–6, 2015.

[16] S. Liu. The color calibration across multi-projector display. *J. of Signal Inf. Process.*, 02:53–58, 2011.

[17] A. Majumder. In *Luminance Management for Seamless Multi-Projector Displays*, volume 36, pages 1506–1509, 2005.

[18] A. Majumder and R. Stevens. Color nonuniformity in projection-based displays: analysis and solutions. *IEEE Trans. on VCG*, 10(2):177–188, 2004.

[19] Aditi Majumder, Cameron Browne, and Michael S Brown. *Practical multi-projector display design*. AK Peters/CRC Press, 2007.

[20] Belen Masia, Gordon Wetzstein, Piotr Didyk, and Diego Gutierrez. A survey on computational displays: Pushing the boundaries of optics, computation, and perception. *Computers & Graphics*, 37(8):1012 – 1038, 2013.

[21] S. Mihara, D. Iwai, and K. Sato. Artifact reduction in radiometric compensation of projector-camera systems for steep reflectance variations. *IEEE Trans. on CSVT*, 24(9):1631–1638, 2014.

[22] Kenneth S. Miller. On the inverse of the sum of matrices. *Mathematics Magazine*, 54(2):67–72, 1981.

[23] Shree K. Nayar, Harish Peri, Michael D. Grossberg, and Peter N. Belhumeur. A projection system with radiometric compensation for screen imperfections. In *Columbia University Technical Report*, 2003.

[24] Isadore Nimeroff, Joan R. Rosenblatt, and Mary C. Dannemiller. Variability of spectral tristimulus values. *J. Opt. Soc. of Am.*, 52, 06 1962.

[25] Matthew Post, Paul Fieguth, Mohamed A Naiel, Zohreh Azimifar, and Mark Lamm. Fast radiometric compensation for nonlinear projectors. In *Proc. Conf. Vis. and Intell. Syst. (CVIS)*, pages 1–3, 2018.

[26] Matthew Post, Paul Fieguth, Mohamed A. Naiel, Zohreh Azimifar, and Mark Lamm. FRESCO: Fast radiometric egocentric screen compensation. In *Proc. IEEE CVPR Workshops*, 2019.

[27] Joaquim Salvi, Jordi Pags, and Joan Batlle. Pattern codification strategies in structured light systems. *PATTERN RECOGNITION*, 37:827–849, 2004.

[28] Gaurav Sharma, Wencheng Wu, and Edul N. Dalal. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application*, 30(1):21–30, 2005.

[29] J. Tsukamoto, D. Iwai, and K. Kashima. Radiometric compensation for cooperative distributed multi-projection system through 2-DOF distributed control. *IEEE Trans. on VCG*, pages 1221–1229, 2015.

[30] D. Wang, I. Sato, T. Okabe, and Y. Sato. Radiometric compensation in a projector-camera system based properties of human vision system. In *Proc. IEEE CVPR Workshops*, pages 70–70, 2005.

[31] Takenobu Yoshida, Chinatsu Horii, and Kosuke Sato. A virtual color reconstruction system for real heritage with light projection. In *Proc. Int. Conf. on VSMM*, pages 1–7, 2003.