University of Kentucky

UKnowledge

University of Kentucky Master's Theses   Graduate School

2007

A CAMERA-BASED ENERGY RELAXATION FRAMEWORK TO MINIMIZE COLOR ARTIFACTS IN A PROJECTED DISPLAY

Nathaniel Sanders
University of Kentucky, nate@natesanders.com

ABSTRACT OF THESIS

A CAMERA-BASED ENERGY RELAXATION FRAMEWORK TO MINIMIZE
COLOR ARTIFACTS IN A PROJECTED DISPLAY

We introduce a technique to automatically correct color inconsistencies in a display
composed of one or more digital light projectors (DLP). The method is agnostic to the
source of error and can detect and address color problems from a number of sources.
Examples include inter- and intra-projector color differences, display surface mark-
ings, and environmental lighting differences on the display. In contrast to methods
that discover and map all colors into the greatest common color space, we minimize
local color discontinuities to create color seamlessness while remaining tolerant to
significant color error. The technique makes use of a commodity camera and high-
dynamic range sensing to measure color gamuts at many different spatial locations.
A differentiable energy function is defined that combines both a smoothness and data
term. This energy function is globally minimized through the successive application
of projective warps defined using gradient descent. After convergence the warps can
be applied at runtime to minimize color defects in the display. The framework is
demonstrated on displays that suffer from several sources of color error.

KEYWORDS: Computer Vision, Projector-Camera Systems, Immersive Displays,
Color, Camera-based Color Correction

Nathaniel Sanders

April 10, 2007

A CAMERA-BASED ENERGY RELAXATION FRAMEWORK TO MINIMIZE
COLOR ARTIFACTS IN A PROJECTED DISPLAY

By

Nathaniel Sanders

Dr. Christopher Jaynes
Director of Dissertation

Dr. Grzegorz Wasilkowski
Director of Graduate Studies

April 10, 2007

RULES FOR THE USE OF THESIS

Unpublished dissertations submitted for the Master's and Doctor's degrees and deposited in the University of Kentucky Library are as a rule open for inspection, but are to be used only with due regard to the rights of the authors. Bibliographical references may be noted, but quotations or summaries of parts may be published only with the permission of the author, and with the usual scholarly acknowledgments.

Extensive copying or publication of the dissertation in whole or in part requires also the consent of the Dean of the Graduate School of the University of Kentucky.
A library which borrows this dissertation for use by its patrons is expected to secure the signature of each user.

<u>Name</u>   <u>Date</u>

_____

_____

_____

_____

_____

_____

_____

_____

THESIS

Nathaniel Sanders

The Graduate School

University of Kentucky

2007

A CAMERA-BASED ENERGY RELAXATION FRAMEWORK TO MINIMIZE
COLOR ARTIFACTS IN A PROJECTED DISPLAY

THESIS

A dissertation submitted in partial fulfillment of the
requirements of the degree of Master of Science in the
College of Engineering at the University of Kentucky

By

Nathaniel Sanders

Lexington, Kentucky

Director: Dr. Christopher Jaynes, Department of Computer Science

Lexington, Kentucky

2007

ACKNOWLEDGEMENTS

Table of Contents

List of Figures

List of Files

n_sanders_thesis.pdf (8.8 M)

Chapter 1

Introduction

Over the past fifteen years, visualization has undergone a revolution due to advances in projection display technology. Projectors have become more convenient and accessible in nearly every possible way. They have become cheaper and smaller. They consume less power than before and their resolution has increased. Large-format displays composed of several overlapping projectors have become possible because of these advances and allow for resolutions of several million pixels – far greater than the typical desktop display.

With the advantages of projection displays come challenges, however. Projectors were designed to project onto planar surfaces that are perpendicular to the optical axis of the projector, so curved or off-axis surfaces warp the projected imagery. We need to discover the geometry of the surface we are projecting onto and apply the appropriate inverse pre-warp to our framebuffer in order to make the output look correct to the viewer. If we allow for arbitrary overlap of projector frustums, we also must account for increased brightness in regions of overlap. To create a truly flexible system, we would also like to allow for a heterogeneous group of projectors to make up our display. This requires discovering and attenuating parameters related to intensity and chrominance, since projectors from different manufacturers could be quite different. Even if all of the projectors in a display system are the same make and model, manufacturing anomalies or differently aged bulbs/lamps in the projectors can cause inter-projector variation. Intra-projector variation can exist, as well, possibly due to degradation of the display mechanism over time.

Unlike desktop displays, which have tightly-coupled pixels which are geometrically well-aligned and are nearly identical in color and intensity properties by comparison,

large-format projection displays require either physical solutions or software solutions to realize the same uniformity that desktop display systems exhibit. While manually adjusting projector location/pose for alignment purposes and manually reconfiguring projector color/intensity parameters can achieve a visually pleasing display, we desire automatic methods for overcoming alignment and color/intensity problems.

Cameras can provide important feedback about the geometry, color, and intensity of a projector-based display. New algorithms that exploit the presence of a camera in a display environment have removed many of the obstacles to deploying digital light projectors in novel environments and applications. In particular, researchers are now exploring large tiled displays that are geometrically aligned on both flat [8], and curved surfaces [16], projection on real-world objects [17], and even steerable displays than can appear on many surfaces within an environment [14]. These displays have opened the door to ad-hoc deployment of large displays, uncontrolled conditions, and display mobility. At the same time, however, these displays all suffer from potentially the same problem: lack of color uniformity. Multi-projector displays are composed of a number of projectors whose color characteristics can vary greatly and whose ad-hoc projection on surfaces with color and textures leads to an image that can contain significant artifacts. Furthermore, environmental lighting conditions can vary across the display surface leading to color change in the observed image.

This paper describes a framework to address color uniformity that operates on the observed color gamuts at many different spatial locations within an observing camera. In contrast to previous approaches that compute the greatest common color space of a display and then map each pixel into that color space, this framework is intended to achieve local color smoothness across the display in the presence of color discontinuities without globally reducing the color dynamic range across the entire display.

For example, Figure 1.1 shows a display surface that contains a yellow marking

Figure 1.1: A color aberration on a projected display surface. A yellow light source falls on the display surface and significantly corrupts a region on the display. (a) full image. (b) closeup of aberration. The method must detect and correct these aberrations without reducing the color range of the entire image.

caused by focused light (indicated by a green circle). Uncorrected projection onto the display results in discontinuities in the color space at the edge of the surface marking and a color artifact. The framework presented here minimizes these local discontinuities within a global energy minimization framework that also seeks to keep the observable color space of the display close to the full dynamic range of the contributing projectors. Focused aberrations such as this one make up the one of the worst class of aberrations for a human viewing the display due to the abrupt discontinuity in color. This class of aberrations also presents great difficulty to the final result of color-correcting methods which reduce the entire display's color space to a common one – a single aberration, perhaps on the periphery of the display, will cause the entirety of the display to lose dynamic range, drastically in some cases. Our method doesn't suffer from this problem, as it attempts to smooth color spaces only in the region of the aberration, leaving regions farther away with their original, full dynamic range.

Another class of aberrations that is difficult to deal with is that of surface mark-

ings. This is very similar to focused light in the way that a color correction system would attempt to correct the problem. Again, systems that try to find a global common color space will reduce the dynamic range of the entire display perhaps to correct one small defect. Our system, just as with focused light, can correct in the region of aberration and leave the dynamic range intact elsewhere.

Aberrations caused by ambient light are often easier to deal with due to their more subtle nature, which often has a natural smoothness or fall-off. Nevertheless, in situations where the ambient light of a display system cannot be controlled, it is advantageous to be able to correct for it. In a system that seeks a common color space, ceiling lights, for example, might cause the display near the floor to have a smaller dynamic range. Our system segments the display into regions with color spaces that become gradually larger as you move from the aberrations near the top of the display to the floor where there are no aberrations. Thus, the gamuts remain unchanged.

In particular, we introduce an energy function that combines smoothness and data terms into the global energy function. The smoothness term measures how closely a particular gamut resembles its spatial neighbors, while the data term measures the distance of the current gamut to the maximum gamut for that region. Using standard relaxation as the global optimizer, projective transformations are successively applied to existing gamuts to move along the energy surface defined by this global error function. This is accomplished by computing the gradient direction of the error function with respect to the parameters of the projective transformation and then applying gradient descent at each iteration of the relaxation algorithm.

The main contributions of this thesis are threefold:

1. Measuring the non-linear output response (gamma) of a projector using a camera

2. Measuring the color gamut across many regions of a projector using a camera

3. Correcting aberrations in a projection display by finding a set of gamut transformations that minimize an energy function chosen for a specific purpose with no assumed physical model of the elements of the system nor model of the environment

Because our technique for correcting the color in a display relies on linear transformations of color gamuts, we need the output response of the projector to be linear. We model the nonlinear response of each channel (R, G, B) of the projector as a power function with a single parameter, gamma. Normally this would require a radiometer and a user in the loop taking several measurements per channel. This would require the user informing the system that the measurement had been taken and that it is OK to display the next image to measure. Our method is fully automatic, however, and as soon as a measurement is taken with the camera, the next is displayed, allowing for rapid measurement and thus estimation of the gamma parameter of the power function describing the output response non-linearity.

Another aspect of our system that normally would require a radiometer is that of measuring the color gamut of the regions in the display. Taking measurements for the nonlinear response discussed above is tedious because of the need to take several measurements for many input/output pairs. This is multiplied greatly by the need for the operator of the radiometer to accurately place the radiometer's location on the display surface for several input/output pairs. With a digital camera we are able to measure the color gamut of the projector at a granularity of 2x2 pixel regions without any trouble and down to the single pixel in some cases.

Finally, we build upon the aforementioned pieces to correct the color of the display by minimizing a user-supplied energy function. This function is chosen by the user for his or her specific purposes related to correcting the color properties of the projection display. In addition to the energy function, several other parameters in the system can be specified by the user. One of the most important is that of neigh-

borhood size. In contrast to other researchers' methods to correct color problems in displays, ours was intended to be a local solution, as evidenced by our neighborhood parameter. This allows us, among other things, to preserve dynamic range in areas of the display without aberrations while improving the color characteristics of portions of the display. By varying the neighborhood size parameter and with careful choice of energy function, we can model several global approaches to color correction similar to those of other researchers.

Chapter 2

Previous Work

The problems of color in projected displays are well known and several approaches to measurement and correction have already been introduced. Each of these approaches has focused on correcting a different source of error and, as a result, each has approached the problem quite differently.

The need for color correctness in projection displays goes back to the CAVE [3] virtual reality system in 1993. The original CAVE was a walk-in immersive projection display with 3 rear-projection screens for walls and a front-projected floor. The system allowed for both head and hand tracking, controllable sound, and also produced stereo visuals. One projector is used for each wall with manual edge-alignment so there is no overlap of projectors. The implementors of the CAVE speak of the need for not compromising color quality as compared to a workstation display, but did not focus on color due to the many other problems they had to solve in constructing such an ambitious and novel system.

Another early project that mentions color correctness is The Office of the Future [18]. The Office of the Future was an umbrella for several projects that sought to achieve compelling shared telepresence systems using cameras, projectors, and networking. Camera clients employing real-time computer vision techniques dynamically extracted per-pixel depth and reflectance maps of surfaces in the scene. This information forms a model that can be used for a variety of tasks. Changes in the surfaces can be detected, such as movement of an object or occlusion of a surface, etc. The model of the scene can also be sent over a network so that a user in a distant location can see a rendering of the original user's workspace. The model of the surface also allows projecting onto the surface accurately. It is this task in which color correctness comes

into play. Because The Office of the Future used casually-aligned overlapping projectors, spatial transitions between projectors include a region where approximately double (or higher, in the case of more than 2 projectors) brightness occurs. Because of this, alpha masks have to be used on a per-projector basis to attenutate intensity. This approach does not deal with the colors of surfaces in the workspace environment, but it was a solid step forward in color correctness in projection displays.

2.1 Non-camera-based Methods

Although a camera is a preferable means of sensing parameters of a display due to its speed and ease of use, early attempts to understand and correct large format projection displays drew methods from the Color Science community and used special-purpose devices such as colorimeters and spectroradiometers. Using such a device has a speed/quality trade-off in that we require the user to make coarse granularity high-quality measurements. Methods which use one of these special-purpose color-measuring devices instead of a camera are practical only when assuming little or no intra-projector variation in color and intensity and also become tedious if color calibration is frequently required.

Majumder [11] was the first to attempt correcting chrominance in addition to intensity. This work employed per-channel, per-projector color maps. Crosstalk between channels and intra-projector color differences are not handled in this work. Also the use of a spectroradiometer as the measuring device limits the frequency/speed with which calibration can be performed because a human is required to control the device. This work was important as a first attempt and as a precursor to Luminance Attenutation Maps [12] discussed later.

The first comprehensive analysis of inter-projector color balancing in projection displays was by Stone [21, 20]. Several later papers build on this analysis. This method first uses a spectroradiometer or colorimeter to measure the gamut of each

projector in the multi-projector display. A common gamut for the display can be found from the measured gamuts. Stone's method is more complicated than the more straightforward method of simply computing the intersection of all gamuts. Instead, this intersection idea is tempered by colorimetric principles. Sources of intra-projector color variation are also analyzed, but only physical solutions for this problem are provided.

In work that is quite similar to our own, Wallace [24] corrects color problems due to differences among projectors in a multi-projector display. The approach also utilizes projective warps in order to align gamuts in different regions on the display. However, the method operates at the level of an entire projector so that it is feasible to measure the mean color response of each gamut vertex by hand using a colorimeter. Furthermore, all gamuts are aligned to a single base gamut or the greatest common gamut in the display. While the coarse granularity of this method allows computations to be extremely fast, it also means that it cannot effectively treat intra-projector color problems, display surface markings, focused light, or ambient light. We measure and adjust gamuts on regions that can be as small as a camera pixel and smooth color artifacts without enforcing a greatest common gamut constraint.

2.2 Camera-based Methods

2.2.1 Global Methods

Majumder [12] later considered the problem of inter-projector and intra-projector luminance variation using per-pixel attenuation maps. The method uses a camera to measure intensity at each camera pixel and then applies a luminance attenuation map (LAM) to all projector pixels so that they are of uniform brightness. The method is similar to our own in that the technique operates in the camera space and does not necessarily distinguish between different projectors and overlap regions. However, this early work focused on intensity only. Also, this solution is a global method in

9

that it seeks a single color space for the entire display through several color space intersections. Global methods such as this suffer from the problem of loss of dynamic range throughout the entire display when only a small portion of the display had low dynamic range before color correction.

Later work [10] introduced color and followed a similar approach to reducing each color space into one that is globally common. An important extension, however, allows color spaces to differ insofar as their differences are unobservable with respect to a simple but effective model of human perception [9]. The framework here differs in that we minimize a global energy with emphasis on local color smoothness. However, we expect to explore how the work of Majumder can be integrated with this framework by formulating LAM similarity measures as a differentiable energy function.

2.2.2 Local Methods

Correction of projected images that fall on non-uniformly colored or textured surfaces has also been addressed by Nayar [13]. The approach introduces a simplified parametric model of both the projection and image formation process and then recovers the parameters of the model in a calibration stage. Once modeled, incorrect color responses in the camera (i.e. due to unexpected color on the display) can be attenuated. The method works quite well and has been shown to remove a majority of artifacts due to surface markings. Because the emphasis of the work is in detecting and removing problems due to surface markings, errors due to other sources (i.e. color problems in the projector or ambient illumination) are not taken into account. As a result, extending this approach to more complex parametric models involving multiple projectors, ambient illumination, etc.. is perhaps difficult and would require more sophisticated color calibration.

Related to these efforts are techniques that attempt to predict the intensity and color in a camera given the pixel values to be projected by one or more projectors.

These researchers are typically interested in accurate prediction of the scene in order to perform three-dimensional reconstruction [25], detection and removal of shadows [6, 22], [5], and other vision tasks in the presence of changing illumination [7, 19]. Although these techniques typically do not involve modification of the projected colors, they do involve similarities to our work.

Chapter 3

Method

The framework is composed of three different components. First, the system is calibrated geometrically and the radiometric response of each projector is linearized. A second stage partitions the camera image into a set of small regions and measures the observable gamut corresponding to each region by rendering color calibration images through the display projectors. Finally, a relaxation stage minimizes an energy term that involves a smoothness term as well as a term that minimizes the motion of the gamuts.

The framework minimizes color artifacts as measured in the camera. In this way, the framework is agnostic to the actual source of error and is quite general. However, this generality comes with a price. The dynamic range of a typical color camera is typically far lower than that of a digital light projector. To address this, high-dynamic range images are synthesized from a set of camera images at different apertures [4]. These images are able to capture a very large dynamic range that encompasses the projected display. For the work described in the following sections, each image refers to a 64-bit floating point image that is used for processing.

Because the method is completely automatic and does not make use of an absolute measure of color (i.e. with a colorimeter [24]), the framework cannot be used to *correct* the display. Instead, the goal is to minimize the perceptual artifacts due to color inconsistencies. A color correction can be applied to the input image before the relaxer runs. This can be achieved by allowing a user to adjust color parameters.

Each of the framework's components is discussed in the sections below. Following that, issues related to runtime manipulation of framebuffer color values are discussed and an algorithm is introduced.

3.1 Preliminary Display Calibration

An invertible mapping between each projector pixel and a corresponding camera pixel must be known before color is addressed. Recovering this information has been the subject of significant previous work and is not the focus here [23, 16, 26]. Although the framework will support a variety of display surface shapes, we assume a flat surface and recover relative geometry using homography trees [2]. Radial distortion is also estimated at this stage using Jean-Yves Bouguet's Camera Calibration Toolbox for Matlab [1].

In addition to requiring geometric information, our approach assumes that different color gamuts in the display can be brought into alignment through projective warps. Because the output intensity of each channel of a typical projector is a nonlinear function of the value in the framebuffer, this nonlinearity must be estimated and removed before color processing. For the results here, we assume that the response of channel C can be modeled as a standard gamma function that maps input intensities between 0..255 to output radiance r_c of the form:

$$r_c = I_c{}^{\gamma_c}$$

(3.1)

This is only an approximation of the more complex nonlinear nature of the image formation and projection process but serves to model the most significant source of nonlinearity in each projector. For example, in a typical DLP projector additional nonlinearities can be introduced at high intensities. Of course, without a radiometer, we cannot directly access the corresponding intensity and radiance pair, although we have turned off "white enhancement" in the DLP projector. Instead, we capture

correspondences as follows. Each projector is divided into distinct regions, R_1 and R_2. Region R_2 is filled with a fixed intensity I, while region R_1 is then filled with a dithering pattern where 50% of its pixels are of some intensity $r = I_k$ and 50% of its pixels are intensity level 0. The camera then measures the average observed intensity in two corresponding camera regions C_1 and C_2. This process is repeated while changing the intensity values in R_2 for all intensity values ranging from 0-255. The closest average intensities from the two regions are selected as a correspondence pair (r, I). This process is again repeated by changing the intensity value in R_2 until a sufficient number of correspondences have been measured. For the results shown here, 20 corresponding samples for each channel were measured.

Equation 3.1 is then fit to these corresponding measurements to recover a per-channel gamma value in each projector of the display. This gamma response is removed in subsequent stages by inverting Equation 3.1 in the graphics hardware. The drivers for our nVidia graphics cards allow us to invert the gamma curve so that we have a linearized response.

3.2 Measuring the Gamut

The camera image is first partitioned into a set of N regions, labeled P_1, , P_i ,P_N. Regions can range in size from a single pixel to 1/4 of the image size. Color relaxation operates on an energy function that involves gamut differences between regions and the choice of region size will impact the results. Region size should be small enough to detect and correct the expected size of color artifacts on the display but should be balanced against the computational cost of additional regions. For the results shown here regions of size 16 camera pixels were typically used.

For each region P_i, the corresponding gamut, G_i, is computed by displaying a fullscreen constant-color image for each vertex of the RGB color cube. Each of the eight images are iteratively displayed by all projectors in the display. As each image

is displayed, camera pixels in region P_i are averaged to compute a value for a corresponding vertex on the region's gamut. Because the geometric mapping between camera pixels and projector pixels is known, the pixels corresponding to region P_i that do not correspond to any projector can be eliminated from processing and do not contribute to the gamut measurement.
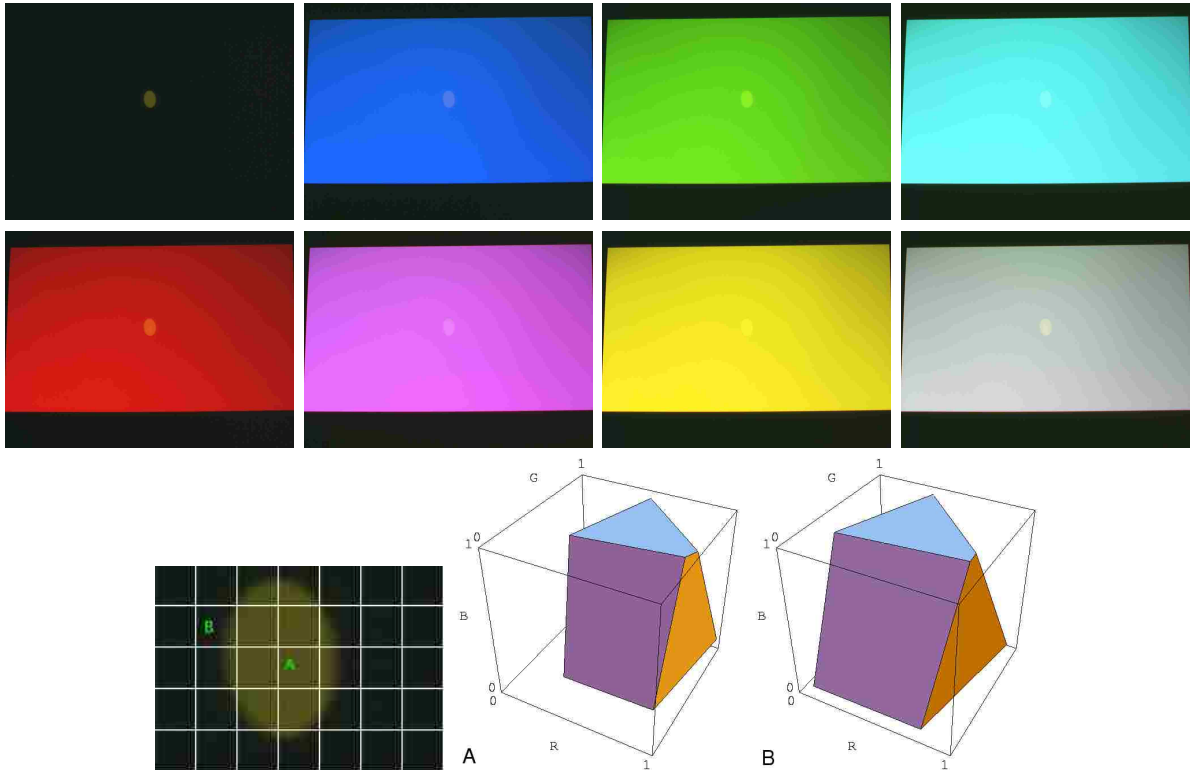
Figure 3.1: Eight input images used to compute the RGB color gamut for each region in the camera space (top two rows). Grid lines are drawn on the aberration to show regions. The measured gamuts corresponding to region A (within the aberration) and region B (partly without the problem region) are also shown. Note that gamut A spans a smaller region on the Red plane (bottom of color space) and has a significantly smaller volume than gamuts outside of the aberration.

Figure 3.1 depicts the eight input images for a display that contains a yellow spot. The measured gamuts corresponding to Regions A and B are shown in Figure 3.1.

Notice that the dynamic range, particularly on the red and green axes of Gamut A is reduced due to the display artifact, while the nearby gamut encompasses a wider range of values. It is these differences that lead to perceptual color artifacts on

the display and, while the projector's dynamic range cannot be increased, the two gamuts should be brought into closer alignment in order to reduce these artifacts. This motivates our choice of energy function described in the next section.

3.3   A Gamut Difference Energy Function

Our definition of a gamut energy function is based on two observations. First, we hope to remove significant artifacts that arise from neighboring regions in the display that have different color characteristics rather than correct the display towards some absolute color space. Examples of these problems include the observable color boundary between two projectors that have different color characteristics or sharp color edges that result when a surface marking influences one region on the display and not its neighbors. Secondly, although there is a theoretically "correct" color space, projector gamuts may only partially overlap that space. Approaches that intersect the gamuts with this global color space can be overly restrictive in many cases (such as dramatic surface markings) and result in a very-low dynamic range display. Instead, our energy measure emphasizes local differences in an attempt to minimize artifacts even in the most challenging conditions.

While we propose a particular form of the energy function here, other energy functions can be used. In fact, we encourage further exploration of energy terms that are more directly inspired by studies of human perception [9], redefine neighbor relations to include inter-reflection artifacts, or utilize a different distance metric between gamuts.

Our energy function is a combination of a data term and a smoothness term. These two terms balance the energy that results from deforming a gamut from its starting shape (the data term) with the energy resulting from the distance of that gamut from the mean of its neighbors (the smoothness term).

The data term energy between the measured gamut G_k and a deformed (proposal)

gamut ΔG, for region k, is given by:

$$D(G_k, \Delta G_k) = \sum_{v_i \in G} \|v_i - \Delta v_i\|^2$$

(3.2)

where v_i represents the RGB 3-vector for gamut vertex i and Δv_i is the corresponding vertex on the proposal gamut.

The smoothness term is defined using a neighborhood relation N over camera regions. For the work here, neighborhoods are defined as the eight-connected regions within a region distance B from the region P_i. All regions within this distance are used to compute a mean gamut for that region k.

$$G_M{}^k = \frac{\sum_{P_n \in N(P_k)} \lambda_n G_n}{\sum_n \lambda_n}, \lambda_k = \left(\frac{1}{\sqrt{2\pi\sigma^2}} e^{\frac{-(x_n - x_k)}{2\sigma^2}}\right)$$

(3.3)

Here $N(P_k)$ is the neighborhood relation on region P_k, G_n is the gamut corresponding to region P_n, and λ_n is the combined weighting coefficient that takes into account the distance between region centers x_n and x_k.

The smoothness energy for proposal gamut, ΔG, is simply the distance of the proposal to its neighborhood mean. This smoothness term is combined with the data into the energy function that describes the error for a particular proposal gamut,

$$E(\Delta G_k) = (1 - \alpha)D(G^k{}_M, \Delta G_k) + \alpha D(G_k, \Delta G_k)$$

(3.4)

where α defines the contribution of the data and smoothness terms to the overall energy function. For the results shown here, alpha values ranging from 0.05 to 0.6 were used.

3.4   Relaxation

Given an energy function that corresponds to the amount of error contained in a particular region, one must define how different gamuts can be proposed in order to minimize this error. For the work here, we utilize a projective transform that maps the initial gamut for region k into a proposal gamut for that region, $G_k = \mathbf{T}G_k$. The projective transform allows significant distortion between gamuts while limiting the number of parameters to 15 that ultimately have to be discovered during the relaxation phase. Projective distortion and similar parametric warps have been used in projector-camera color alignments [24] as well as camera color calibration between multiple cameras.

The choice of both $E(G_k)$ and $\mathbf{T}$ define the energy minimization problem that must be solved. It should be noted that the framework supports different definitions of both of these components and further exploration of these terms is an interesting topic of research.

For notational purposes, we define $G_k(t)$ to be the gamut corresponding to region k at iteration t. The algorithm takes into account three different gamuts for region k at each iteration of the relaxation: the measured gamut $G_k(0)$, the current proposed gamut at iteration i, $G_k(i)$, and the new proposed gamut, $T_{i+1} \, G_k(i)$, warped by a new transform, T_{i+1}.

The method begins by computing the energy function for each region in the display (see Section 3.3) and sorting the list of regions in descending order of energy. The top

element of this list is selected and a new proposal transform is generated in a two-part process. First we compute an approximate target gamut (Equation 3.5) by taking a linear combination of the mean gamut of the neighborhood and the measured gamut. The weights are $(1-\alpha)$ and α, as discussed above. We then compute, via Equation 3.6 the mean of the translation vectors from corresponding vertices between the current gamut and the approximate target gamut. We translate the current gamut by this mean translation vector. This gamut-alignment step is beneficial because simply using gradient descent without using the mean-vector-translation step is more likely to encounter local minima.

$$G_{approx} = (1\text{-}\alpha)\mathrm{G}_M^k + \alpha\mathrm{G}_k$$

(3.5)

$$\vec{t} = (1/8)\sum_{j=1}^{8}(G_{approx}[j] - G_k(i)[j])$$

(3.6)

where $G[j]$ is the homogeneous point corresponding to the jth vertex of gamut G. Now we need to update our current transform $\mathbf{T}(\mathrm{i})$ by the mean translation we have computed. We do this simply by copying the values from $\vec{t}$ into the rightmost column of a 4x4 identity matrix, so that we have a pure translation matrix $\mathbf{T}_{trans}$. Our new current matrix is $\mathbf{T}_{trans}\mathbf{T}(i)$ and our current putative gamut is $G_{intermediate} = \mathbf{T}_{trans}\mathbf{T}(i)G_k(0)$.

After the relatively large translation, we perform local optimization by first computing the maximum gradient descent defined by the Jacobian of the energy term with respect to the parameters of the transformation:

$$\vec{p} = \frac{\partial E(G_{intermediate})}{\partial \mathbf{T}}$$

(3.7)

Given our choice of $\mathbf{T}$, $\vec{p}$ is a 15-dimensional vector describing the direction of maximum descent. A line search is performed along $\vec{p}$. We solve our line search problem by using a technique called Brent's Method [15]. The solution to the line search problem is called δ. The minimum value of the energy function is used to determine the new gamut transform for that region, $\mathbf{T}(i+1) = \mathbf{T}_{trans}\mathbf{T}(i) + \delta *$ $unstack(-\vec{p})$, that maps the region's gamut to an updated gamut with less energy.

If the new transformation does not, in fact, reduce the energy for the whole neighborhood, we back out all changes for this iteration and the entire process is applied to the next region in the sorted list. Note that our goal is to reduce the entire system's energy by operating on its current worst member. Checking to see that the entire system's energy is reduced is equivalent to checking that the neighborhood's energy has been reduced because changing a gamut can only affect the energy of regions within its neighborhood – no regions outside the neighborhood can possibly be affected. If the energy measure for any region is reduced, the energy for all regions that could be impacted by this change (i.e. all regions in the neighborhood) is recomputed and a new sorted list of gamut regions is computed.

The algorithm then iterates in this manner on each new gamut list as it is produced until a stopping criterion is reached. For the results here, two stopping criteria are used. First, if the total energy of the error surface falls below some threshold, the

relaxation algorithm terminates. The algorithm will also terminate if the energy difference between the current gamut and the resulting gamut for any region visited on the current iteration is below a threshold. Figure 3.2 depicts this iterative process on the input data from Figure 3.1. The number of times a region is visited and modified is encoded as intensity in Figure 3.2a.

(a)   (b)
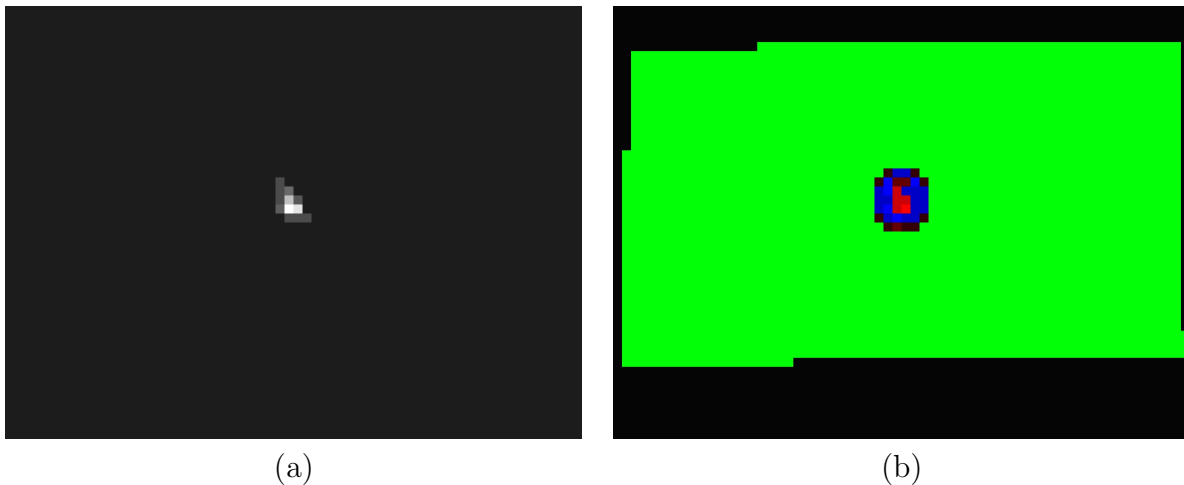
Figure 3.2: Visualization of the energy relaxation behavior for input images in Figure 3.1. Gamut relaxation operates primarily on discontinuities in the energy surface. (a) Intensity encodes the number of times a particular region is visited and modified in the relaxation process. (b) Volume increase is indicated by red, decrease by blue, no change by green.

Notice that regions on the border of the color artifact are visited more often as they are corrected, impact the overall energy, and then are revisited. More distant regions are corrected far less. Figure 3.2b encodes the volume change between the starting and optimized gamuts. Gamuts on the border of the artifact were reduced in size to better match their neighbors on the interior of the artifact where the color dynamic range of the display is reduced.

The choice of particular stopping criterion is related to the error function used as well as the choice of transformation. For example, if one defines an energy function based on traits of the human perceptual system, a stopping criterion may be defined based on the minimum observable change in color gamuts.

3.5 Correction

(a) (b) (c)
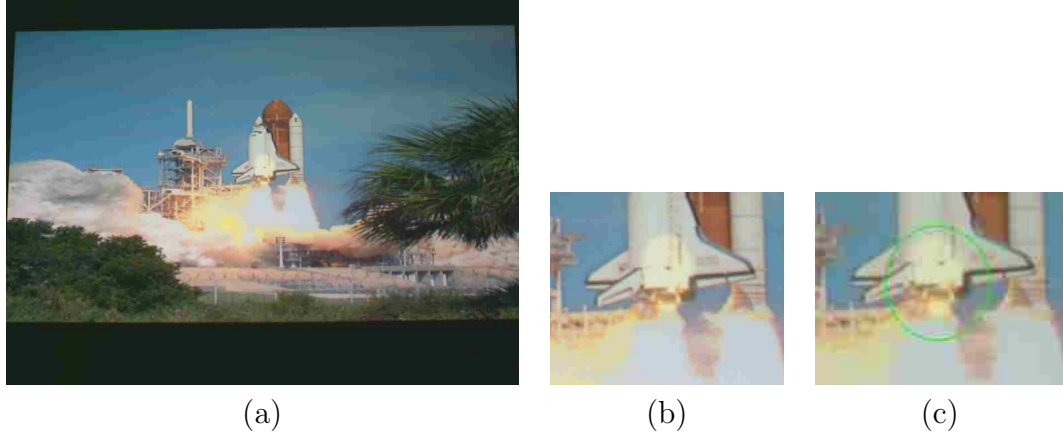
Figure 3.3: Corrected image for the example input image from Figures 1.1- 3.2. (a) Result of correction as seen from the camera. Because the technique imposes local smoothness constraints without requiring the entire image to lie within a common color space. The result image still exhibits good color contrast in regions far from the aberration. (b) Close-up of corrected areas near the aberration. The yellow color aberration is less saturated. More importantly, the edge between this region and its neighbors is less apparent than the uncorrected image. (c) Close-up of uncorrected image shown for comparison.

After relaxation, we have the final transformations in each region and can apply them to pixels in the input image. Because we have committed to a space of linear color space transforms, the resulting transformations can be applied simply by multiplying incoming color values with the transformation corresponding to the region in which the pixel lies. For example, given a pixel to be displayed in region k, with color value of c, a new color value is computed as:

$$\hat{c}_k = \mathbf{T}[r \ g \ b \ 1]^T$$

(3.8)

After transformation, the output color values are clamped to the range 0-255 and

displayed. Figure 3.3 shows the corrected and uncorrected images from the example shown in Figures 3.1 and 3.2.

Chapter 4

Experimental Results

4.1 System Setup and Composition

The image in Figure 4.1 shows a typical experimental setup. The equipment used to build the system was

1. Hitachi HV-F31 3-CCD IEEE-1394 camera mounted on tripod transmitting 1024x768 24-bit color RGB still images

2. 2 Optoma H57 DLP projectors with native resolution of 1024x576; DVI input interface; white enhancement turned off; 1 projector used for focused-light aberrations, 1 used for image display

3. Camera server: Pentium 4 2.8GHz 512 MB RAM, Ethernet, 100GB hard drive

4. Image Projector server: Pentium 4 3.06GHz, 1GB RAM, Ethernet, 100GB hard drive – all computational operations are performed on this machine

5. Aberration Projector server: Pentium 4 2.8GHz 512 MB RAM, Ethernet, 100GB hard drive

6. Lamp with a variety of colored bulbs for ambient-light aberrations

The majority of the system is implemented in C++ with OpenGL for 2D rendering and Intel's OpenCV library for matrix math and image loading and saving. The gradient direction computations are calculated using the GNU Scientific Library and the gradient length computations are performed using the Brent's method code from *Numerical Recipes in C* [15]. The solver for high dynamic range imaging is a port of

Figure 4.1: A projector-camera system equipped for image display as well as focused-light and ambient-light aberrations.

Paul Debevec's Matlab code to the GNU Octave language. Several utility programs for managing the large datasets involved in this project, parsing log files, and generating visualizations of the data were written in combinations of C++, Octave, Perl, shell scripts, and the ImageMagick command-line utilities. The camera server runs on Windows and is controlled by using the National Instruments NI-IMAQ API.

4.2 Evaluation

We tested the system for its ability to increase the similarity between gamuts in an area where an aberration occurred. Figure 4.2(a) shows the black calibration image displaying a focused-light multi-color curved-line artifact. Figure 4.2(b) shows the input image we wish to correct.

Selecting a method to interpret the results is not trivial. Attempting to apply seemingly obvious or standard ways of analyzing results such as SAD/SSD or perhaps measuring the maximum gradient of the projected image near the aberration

is misleading. This is because doing so for a particular parameter set would inadvertently evaluate the system operator's ability to select parameters for a desired outcome. Instead, we report the percentage of energy decrease across the display and also take a look at some compelling visualizations of various aspects of the system and its results.

The system was run to measure initial gamuts and then correct the input image. In this experiment we used a neighborhood size of 15x15 regions and $\alpha = 0.05$. Each region was two pixels by two pixels. After 1000 iterations of the relaxer (near convergence), we compared the initial energy to the final energy. The energy over the entire region space in this experiment decreased from 68.2416 units to 36.879 units for 45.96% decrease in energy. This demonstrates that the technique correctly adjusts gamuts towards a lower energy state based on the definition of the energy function.

In Figure 4.3(a) and (b) we see the uncorrected and corrected versions of the images, respectively. This aberration is quite strong in that it cuts a sharp edge through the display, so it is necessary to zoom in to see where the correction happens. Figure 4.3(c) and (d) show a close-up of the blue portion of the aberration. Here we can more easily see the softening of the blue as it it appears more strongly in the neighborhood around the blue line to soften the edge. It is possible in this case that if the purpose of this parameter set was to smooth the aberration, a lower value of α should have been used to weigh the mean gamut more heavily and the measured gamut less heavily. We may also have benefited from a larger neighborhood size, which would make the area of effect around the aberration larger.

In Figure 4.4(a), we show the difference image between the camera images of the two aberrated images (corrected and uncorrected). In this image, the intensities of the difference have been multiplied by five for emphasis. Here it can be seen that the correction forms a "shell" around the aberration which results in a smooth blend from the aberration to the background. Another visualization that has strong similarities

to the aforementioned difference image is Figure 4.4(b), however while Figure 4.4(a) was computed as the pixel-difference between two camera images, Figure 4.4(b) is a visualization of the black-point translation from the measured (initial) gamuts in each region to the gamuts after near-convergence (1000 iterations, in this case). In other words, for each region, we subtracted the initial black point from the final black point and plotted that 3-vector as RGB values. Gamut movement, in general, is difficult to visualize because of its 3D nature. In order to more easily understand how the gamuts have been transformed, we look at black-point-translation vectors or volume scalars to enable us to make useful inferences about the nature of the data. In this particular case, there is significant black-point translation, which agrees with intuition, as our focused-light aberration should, for the most part, simply be additive energy. Inside the boxed area in Figure 4.4(c), we see that the shell formed around the aberration is of high intensity near the aberration and decreases with distance.

The last of the visualizations shows the volume in each region both at the time of measurement and after 1000 iterations. Figure 4.5 shows the "before" and "after" images, respectively, while the image to the right of each larger image shows a close-up of the volume visualization in the region of aberration. In Figure 4.5(d), we can see a medium-grey-colored region in between the darkly-colored aberration portion and the large, lighter-grey-colored unaberrated region. This is where the relaxer was taking "normal" gamuts and making them more like the gamuts in the aberrated region, in accordance with this particular parameter set which decreases energy based on similarity to a gaussian-weighted mean of neighbors. The bright white specks apparent in some parts of the transformed area are regions that have grown in volume. This is an anomalous case that occurs due to local minima.

Table 4.1 shows some values for percent decrease in energy for 3 experiments with all parameters equal except for α. These experiments used the same aberration and thus the same input calibration images as the previous experiment. These three

experiments show the effect of α on the system. With all other parameters equal, as alpha gets smaller, percent decrease in energy goes up. This is to be expected, as gamuts are allowed to change more the smaller α is. If α were 1.00, for example, gamuts would would achieve their lowest energy state by staying at measured value and we would have no energy decrease.

Region size	α	Iterations	Neighborhood size	% Energy decrease
4	0.05	1000	9x9	47.86%
4	0.2	1000	9x9	45.74%
4	0.6	1000	9x9	39.07%

Table 4.1: Percent energy decrease for a given parameter set

4.3 Limitations

We have demonstrated a system with a solid theoretical framework. However, during the course of experimentation, we have discovered limitations in our work. The volume visualizations in Figure 4.5 show that our relaxer allows some gamuts to grow, which is physically impossible. Also, the amount of correction in Figure 4.3 is so low as to be difficult to see even when magnified. This is likely due to a sub-optimal choice of input parameters. More experimentation and better understanding of how input parameters affect the course of the relaxer will lead to better results.

(a)

(b)

Figure 4.2: (a) The focused-light aberration for this experiment. (b) The unaberrated input image.

(a)   (b)

(c)   (d)

Figure 4.3: Uncorrected and corrected camera images with close-up of aberration portion. (a)The aberrated/*uncorrected* version of the input image. (b) The aberrated/*corrected* version of the input image. (c) Blue portion of aberration region of *uncorrected* image magnified. (d) Blue portion of aberration region of *corrected* image magnified.

(a)	(b)	(c)

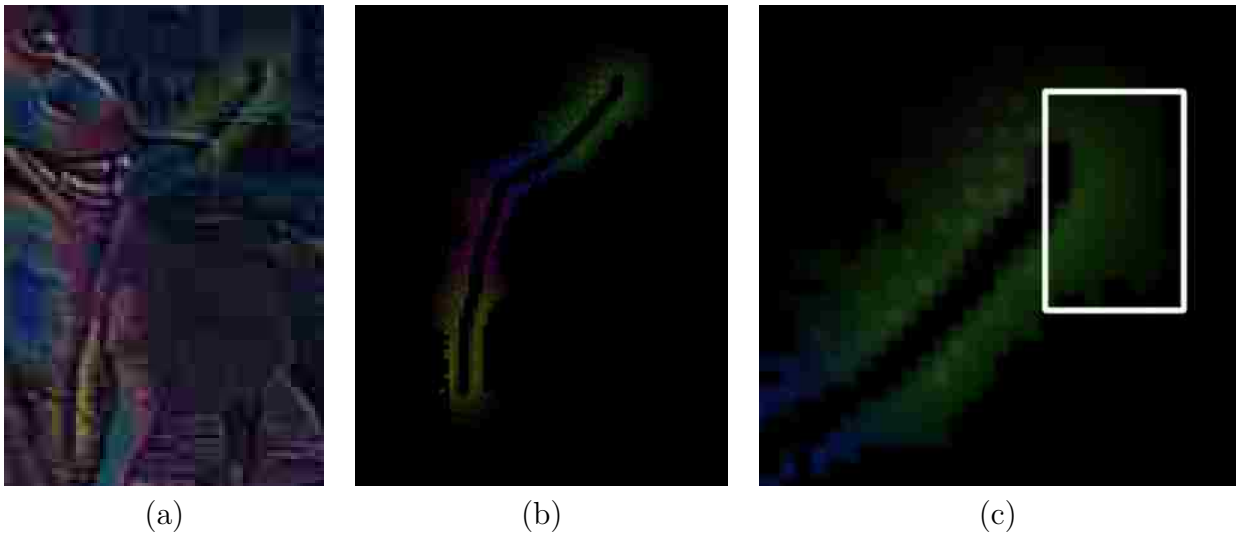Figure 4.4: Some visualizations of the process: (a)The difference image (multiplied by 5 for visualization purposes) between the uncorrected and corrected images in the region of the aberration. (b) magnification of the black-point translation image (values multiplied by 5 for visualization purposes) (c) further zoom of black-point translation image with box emphasizing correction fall-off/smoothness
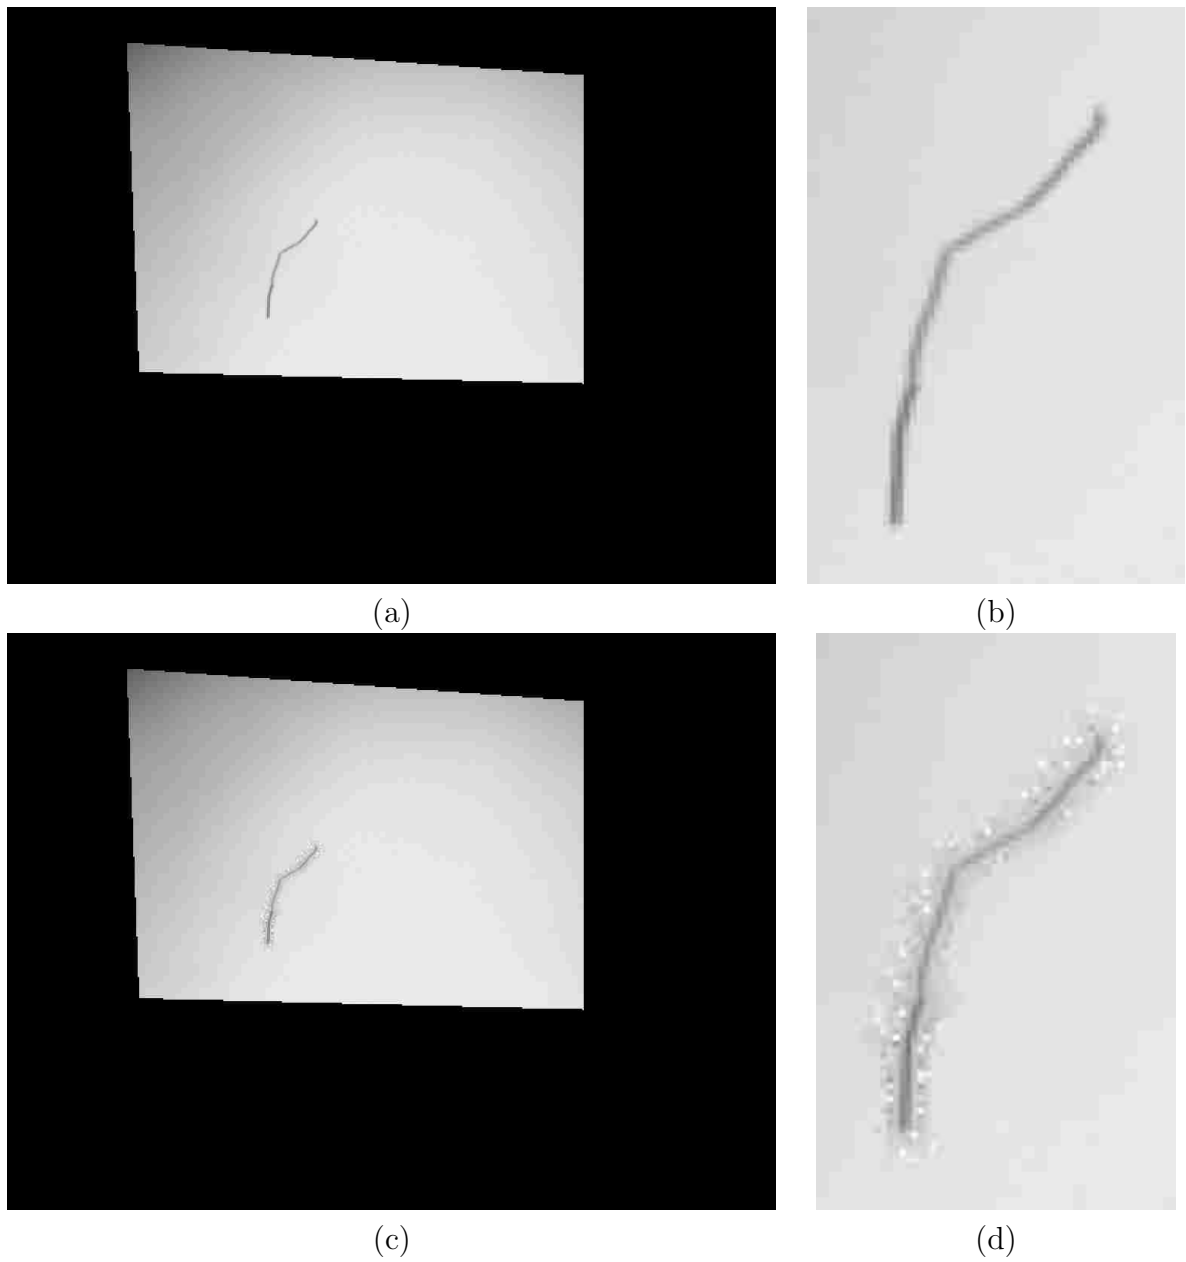
(a)

(b)

(c)

(d)

Figure 4.5: Visualization of volumes of gamuts across all regions (darker color indicates smaller gamut): (a) at measurement time/before correction (b) magnification before correction (c) after correction (d) magnification after correction

Chapter 5

Conclusions and Future Work

We have introduced a general system for transforming color in a projected display environment. To meet this end, we have also contributed two key sub-systems that measure the non-linear output response of a projector and the color gamut across a display by using a commodity camera.

The full extent and flexibility of the system certainly could not be captured in this document alone. The relaxer engine combined with variable region sizes, neighborhood sizes, filters other than a gaussian-weighted mean, and weighting coefficients create a staggering number of inputs and outputs. Future work on this system should include a better understanding of how to select input parameters for typical tasks. For example, how might we select neighborhood size and region size as a function of aberration size? In addition to those more basic questions, future research may attempt to phrase other researchers' algorithms or desired ends in terms of our system. Through weighting coefficients, multiple ends could be balanced in different ways, just as we balanced between the measured gamut (the data term) and the mean of our neighbors (the smoothness term) in the basic case in our research. Also, while the work shown here features relatively small neighborhood sizes, very large neighborhood sizes (perhaps the size of the entire display) could be used to approximate algorithms which take into account global information. As an example of this, a very large neighborhood coupled with an order-statistic filter which selects the gamut with least volume may closely approximate one of the "greatest common gamut"-style algorithms which have been discussed by other researchers. The system described here could provide a test-bed for the creation of new ideas, as well as the hybridization/combination of existing ideas for color transformation in projection displays.

Bibliography

[1] J.-Y. Bouguet, *Camera calibration toolbox for matlab*. http://www.vision.caltech.edu/bouguetj/calib_doc/.

[2] H. Chen, R. Sukthankar, G. Wallace, and T. Cham, *Calibrating scalable multi-projector displays using camera homography trees*, 2001.

[3] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti, *Surround-screen projection-based virtual reality: the design and implementation of the cave*, in SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques, New York, NY, USA, 1993, ACM Press, pp. 135–142.

[4] P. E. Debevec and J. Malik, *Recovering high dynamic range radiance maps from photographs*, Computer Graphics, 31 (1997), pp. 369–378.

[5] M. N. Hilario and J. R. Cooperstock, *Occlusion detection for front-projected interactive dsiplays*, in Advances in Pervasive Computing, A. Ferscha, H. Hoertner, and G. Kotsis, eds., Vienna, Austria, Apr. 2004, Austrian Computer Society (OCG), pp. 265–271.

[6] C. Jaynes, S. Webb, R. M. Steele, M. Brown, and W. B. Seales, *Dynamic shadow removal from front projection displays*, in VIS '01: Proceedings of the conference on Visualization '01, Washington, DC, USA, 2001, IEEE Computer Society, pp. 175–182.

[7] A. Kale, K. Kwan, and C. Jaynes, *Epipolar contrained user pushbutton selection in projected interfaces*, in CVPRW '04: Proceedings of the 2004 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'04) Volume 10, Washington, DC, USA, 2004, IEEE Computer Society, p. 156.

[8] K. Li, H. Chen, Y. Chen, D. W. Clark, P. Cook, S. Damianakis, G. Essl, A. Finkelstein, T. Funkhouser, T. Housel, A. Klein, Z. Liu, E. Praun, R. Samanta, B. Shedd, J. P. Singh, G. Tzanetakis, and J. Zheng, *Building and using a scalable display wall system*, IEEE Comput. Graph. Appl., 20 (2000), pp. 29–37.

[9] A. Majumder, *Contrast enhancement of multi-displays using human contrast sensitivity*, in CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 2, Washington, DC, USA, 2005, IEEE Computer Society, pp. 377–382.

[10] A. Majumder and M. Gopi, *Modeling color properties of tiled displays*, Computer Graphics Forum, 24 (2005), pp. 149–163.

[11] A. Majumder, Z. He, H. Towles, and G. Welch, *Achieving color uniformity across multi-projector displays*, in VIS '00: Proceedings of the conference

on Visualization '00, Los Alamitos, CA, USA, 2000, IEEE Computer Society Press, pp. 117–124.

[12] A. MAJUMDER AND R. STEVENS, *Lam: luminance attenuation map for photometric uniformity in projection based displays*, in VRST '02: Proceedings of the ACM symposium on Virtual reality software and technology, New York, NY, USA, 2002, ACM Press, pp. 147–154.

[13] S. NAYAR, H. PERI, M. GROSSBERG, AND P. BELHUMEUR, *A Projection System with Radiometric Compensation for Screen Imperfections*, in ICCV Workshop on Projector-Camera Systems (PROCAMS), Oct 2003.

[14] C. PINHANEZ, *Using a steerable projector and a camera to transform surfaces into interactive displays*, in CHI '01: CHI '01 extended abstracts on Human factors in computing systems, New York, NY, USA, 2001, ACM Press, pp. 369–370.

[15] W. H. PRESS, B. P. FLANNERY, S. A. TEUKOLSKY, AND W. T. VETTERLING, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, Cambridge (UK) and New York, 2nd ed., 1992.

[16] R. RASKAR, M. S. BROWN, R. YANG, W.-C. CHEN, G. WELCH, H. TOWLES, W. B. SEALES, AND H. FUCHS, *Multi-projector displays using camera-based registration*, in IEEE Visualization, 1999, pp. 161–168.

[17] R. RASKAR, J. VAN BAAR, P. BEARDSLEY, T. WILLWACHER, S. RAO, AND C. FORLINES, *ilamps: geometrically aware and self-configuring projectors*, ACM Trans. Graph., 22 (2003), pp. 809–818.

[18] R. RASKAR, G. WELCH, M. CUTTS, A. LAKE, L. STESIN, AND H. FUCHS, *The office of the future: a unified approach to image-based modeling and spatially immersive displays*, in SIGGRAPH '98: Proceedings of the 25th annual conference on Computer graphics and interactive techniques, New York, NY, USA, 1998, ACM Press, pp. 179–188.

[19] P. SONG AND T.-J. CHAM, *A theory for photometric self-calibration of multiple overlapping projectors and cameras*, in CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Workshops, Washington, DC, USA, 2005, IEEE Computer Society, p. 97.

[20] M. STONE, *Color balancing experimental projection displays*, in IS&T 9th Color Imaging Conference, vol. 9, 2001.

[21] M. C. STONE, *Color and brightness appearance issues for tiled displays*, IEEE Computer Graphics and Applications, 21 (2001), pp. 58–66.

[22] R. SUKTHANKAR, T. CHAM, AND G. SUKTHANKAR, *Dynamic shadow elimination for multi-projector displays*, 2001.

[23] R. SURATI, *Scalable self-calibrating display technology for seamless large-scale displays*, 1999.

[24] G. WALLACE, H. CHEN, AND K. LI, *Color gamut matching for tiled display walls*, in EGVE '03: Proceedings of the workshop on Virtual environments 2003, New York, NY, USA, 2003, ACM Press, pp. 293–302.

[25] G. WELCH, R. YANG, M. BRUCE CAIRNS, H. TOWLES, A. STATE, A. ILIE, S. BECKER, D. RUSSO, J. FUNARO, D. SONNENWALD, K. MAYER-PATEL, B. D. ALLEN, H. YANG, M. EUGENE FREID, A. VAN DAM, , AND H. FUCHS, *3d telepresence for off-line surgical training and on-line remote consultation*, in Proceedings of ICAT CREST Symposium on Telecommunication, Teleimmersion, and Telexistence, 2004.

[26] R. YANG, D. GOTZ, J. HENSLEY, H. TOWLES, AND M. BROWN, *Pixelflex: A reconfigurable multi-projector display system*, 2001.

Vita

Nathaniel Sanders

Background

Date of Birth: December 19, 1975

Place of Birth: Louisville, Kentucky

Education

Bachelor of Science in Computer Science, University of Kentucky, 1999.

Professional Experience

May 1997 - August 1998. Intern/Developer at Perot Systems Corporation. Lexington, KY.

August 1998 - June 1999. Developer at CommerceInc. Lexington, KY.

2000-2001. Independent Consultant/Software Developer. Lexington, KY.

February 2002 - August 2003. System Administrator/Research Programmer at the Laboratory for Advanced Networking and later the Center for Visualization and Virtual Environments. University of Kentucky, Lexington, KY.

August 2003 - May 2007. Research Assistant at Center for Visualization and Virtual Environments. University of Kentucky, Lexington, KY.

Publications

C. Jaynes, A. Kale, N. Sanders, and E. Grossman. *The Terrascope Dataset: A Scripted Multi-Camera Indoor Video Surveillance Dataset with Ground-truth*, Proceedings of the IEEE Workshop on VS PETS, 2005.

N. Sanders and C. Jaynes. *Class-specific Color Camera Calibration with Application to Object Recognition*, WACV 2005, 2005.

N. Sanders and C. Jaynes. *A Camera-Based Energy Relaxation Framework to Minimize Color Artifacts in a Projected Display*, 2006 Conference on Computer Vision and Pattern Recognition Workshop, 2006.