

The Wall of Inspiration: A Computer Aided Color Selection System

Seth Berrier^{1,2}, Clement Shimizu^{1,2}, Patrick Chong³, D'nardo Colucci², Gary Meyer^{1,2}

¹ University of Minnesota, Department of Computer Science and Engineering; Minneapolis, Minnesota

² University of Minnesota, Digital Technology Center; Minneapolis, Minnesota

³ Benjamin Moore & Co.; Flanders, New Jersey

Abstract

Computer graphics hardware and software is used to create a new set of tools for selecting a household paint color. Important surface reflection effects, such as gloss and pearlescence, are rendered in real time using pixel shading hardware. Unlike traditional paint chips, this feature permits colors to be evaluated on three dimensional shapes and in different surface finishes. A novel set of color navigation tools is provided to allow the user to visualize the color collection supplied by the manufacturer, to select subsets of the collection for consideration, and to match or modify an existing color sample. This interface is easier to use and less expensive to produce than the fan deck found in most paint stores. The pictures produced by the software are displayed on a four panel video wall. The large size of the wall allows the user to evaluate the appearance of the color at a scale similar to its final dimension once it has been applied to a surface.

Introduction

Computer graphic simulation has the potential to play an important role in the marketing of paint and other surface coatings. Digital control of additive color reproduction devices makes it possible to achieve very accurate color rendition. Increasing video projector brightness and resolution allows dramatic video walls to be considered for use in paint stores and home improvement centers. The recent introduction of per pixel shading hardware permits complex color appearance calculations to be performed and displayed in real time. Taken together these advancements make possible new methods of using computer graphics to present virtual paint samples to consumers. In this paper we describe the development of the Wall of Inspiration: the first retail display in the paint industry to take full advantage of computer graphic simulation based on the use of modern programmable shader hardware as well as large video wall technology.

We begin by considering the paint selection tools that are currently available. The first of these aids are small chips of color made from non-paint materials that very closely resemble the color of the finished product. These chips measure anywhere from a few centimeters to half a meter in size depending on the manufacturer. The larger ones may cost money whereas the smaller samples are usually free. The two most common forms of these paint chips are strips of colors found in a paint store display rack or decks that fan out around a single retaining pin. More recently, paint manufacturers have started offering small jars or pouches of pre-mixed paint that will cover a piece of wall roughly 3 sq. ft. in size. These samples are actual latex-based paint and are intended to be viewed on the wall they would eventually cover. Unlike the small chips, these samples cost a few dollars. Even more recently, small pieces of wallpaper (roughly

5 sq. ft. in size) have appeared that approximate the appearance of the paint and cost about the same as the sample jars.

Paint manufacturers try to minimize their cost for making these color selection tools. When the expense gets too high (as in the case of the sample jars and the wallpaper) it must be passed on to the consumer. Considering the large quantities of physical color samples that are produced, more versatile solutions with a potentially large upfront investment but little or no manufacturing expense are desirable. To this end, some manufacturers have developed special software for the paint selection task. They make it freely available on the Internet or on computers installed in the paint stores. Some even offer to sell the software to the consumer so that it can be used on their home PC. One such program is *Personal Color Viewer*[®] (see Figure 1). This software allows the user to take a photograph of the wall or room to be painted and replace the existing color with Benjamin Moore[®] paint colors. While the technique is straightforward (no geometry is present, the rendering is entirely done with a single photograph) the end result can be convincing.

Each of these tools has drawbacks. The small color chips are too small to give an effective impression of the final product. In their most common form (the strips in the display rack) several colors are presented at once in close proximity to one another. This allows simultaneous contrast to skew the colors slightly. The fan deck is restrictive in that the colors are locked together in a specific order. It is impossible to remove a single strip or a single color so it can be viewed in context. Also, the chips are necessarily flat making it difficult to assess the quality of gloss or of special metallic or pearlescent paints. Both the

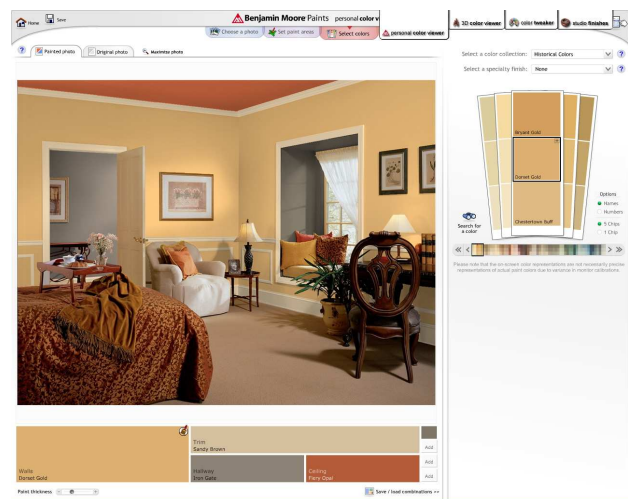


Figure 1. Benjamin Moore's Personal Color Viewer[®] software



Figure 2. The Room of Inspiration

fan deck and the display rack are organized in a somewhat non-intuitive manner. Individual collections of colors are kept separate even though they contain very similar hues. This creates the problem that colors that are perceptually close to one another are not physically close to one another in the display rack.

The sample jars of paint and the colored wallpaper sections are larger and can be used in context but are still far from the scale of the entire wall. As mentioned previously, the higher manufacturing cost of these color selection tools requires the consumer to bear some of the financial burden. The success of the final rendering from the Personal Color Viewer[®] software is very dependent on the quality of the original photograph. The Personal Color Viewer[®] software comes with a large collection of beautiful, studio-quality photographs of professionally decorated rooms that produce wonderful end results. However, when the same techniques are used on personal photographs taken with low quality cameras of rooms with poor lighting and amateur decorating, the end results can be disappointing.

In an effort to alleviate these shortcomings, we have created a new paint selection tool using modern techniques in computer graphics and current display hardware. Our system has been implemented and installed in a New York City paint store. The system, called *The Wall of Inspiration*, has two major components: 1) a large scale rear projected video display driven by graphic cards on personal computers, and 2) custom software that runs on the wall and provides the consumer with unique tools to select paint colors (see Figure 2). The large scale is intended to overcome the size problem of the small chips. The custom software incorporates features designed specifically to address problems presented by the current color selection tools.

Unlike current tools, our software accurately simulates different gloss levels and metallic reflection properties using existing techniques in computer graphics. These special appearance properties are rendered on three-dimensional objects illuminated by photographically captured environments. These environments range from a room with professionally designed studio lighting to large day-lit outdoor spaces to the exact room in which the subject is using the software. To evaluate the appearance of the paint we have provided a representative set of three-dimensional objects including a piece of crown molding, a door, a wooden chair, a ceiling medallion and even a radiator. We harness modern graphics hardware to generate the appearance effects and manipulate the objects in real-time at fully interactive speeds. Our system has different environments available for reflection calculations as well as different sets of fine-tuned RGB values that

have been selected for a particular monitor and illuminant.

In this paper we will first discuss the details of the software system focusing on the user experience. It is shown how each part of the system attempts to overcome the problems of current color selection aids. The next section describes the software implementation by listing the hardware used and the computer graphics techniques employed. The rendering engine and the graphics algorithms utilized in simulating the appearance of the paint are covered in this section. Finally the work is summarized and some conclusions are drawn.

Our System

Our software is designed around three color selection tasks:

- Visualizing paint colors on complex 3D objects
- Visualizing and modifying a color that has been acquired with a spectrophotometer
- Navigating a large collection of colors

The 3D Color Viewer

To perceive the detailed reflective properties of a paint it is helpful to view it on an object with complex geometry. This is particularly the case when evaluating metallic or glossy reflections. A module called *The 3D Color Viewer* was created in our system for this task (see Figure 3). This module provides a library of common household objects. Any paint color from our database of pre-sampled colors can be placed on the surface of these objects. In addition, all of these paints are available in several gloss levels ranging from matte to high-gloss. The 3D Color Viewer is able to simulate five distinct gloss levels for each paint color (see Figure 4). This module is an attempt to overcome some of the problems inherent in the small chips of color that are available from the manufacturer which are inherently flat and lacking detailed geometry and which are only available in a single gloss finish.

Our system also has the ability to visualize certain goniochromatic colors; particularly metallic and pearlescent tints which are becoming increasingly more popular in consumer interior design applications. These paints possess special reflective properties that cause their hue and chroma to shift slightly as the viewing angle changes. It is very difficult to see this shift on a flat chip lacking detailed geometry. By placing these colors on complex objects with curvature and fine detail, the unique reflective properties of these colors are much easier to perceive than when




Figure 3. The 3D Color Viewer Module




Figure 4. Gloss levels - matte, eggshell, pearl, semi-gloss & high gloss

viewed on a flat paint chip. The ceiling medallion in Figure 3 is rendered with a copper version of this paint.

The 3D Color Viewer also allows the user to grab the object and rotate it in real time. This enables them to explore the full range of reflectance angles and lighting across the surface. Moving the object in front of a light source is vital for evaluating glossy finishes and metallic paints.

The Color Tweaker

Often times a customer may bring in a swatch of cloth or a similar object that has the color they are interested in. However, they do not want that exact color. For example, they may want a “happier” color or a more “mellow” color. This task was identified as tweaking a color. To accommodate this, another module was developed called *The Color Tweaker* (see Figure 5). The Color Tweaker is a decentralized visualization of the entire database of colors in CIE $L^*a^*b^*$ color space. Each color is rendered as a small sphere centered at its $L^*a^*b^*$ value (calculated in a pre-processing step and available in the database). The camera is set up to look directly at the coordinates of the customer’s color (determined by scanning it with an in-store spectrophotometer). A set of axes are also drawn that originate from the point of the scanned color and identify the lightness, chroma and hue dimensions as defined by the $L^*a^*b^*$ version of CIE LCH space.

Integrated into the axes are several beads that can be used to limit the number of visible colors along each of these dimensions. Moving the beads with the mouse controls a visible wedge shaped volume extent. Only colors that fall within this volume are drawn, all others are clipped. Since the hue axis is curved in $L^*a^*b^*$ space, the shape that this extent creates is actually a segment of a washer with some thickness. The volume extent

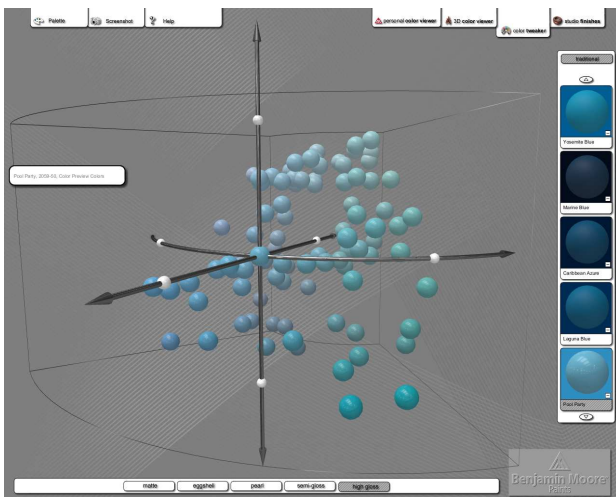


Figure 5. The Color Tweaker Module

and the viewing direction stays centered at the location of the scanned color until the user double-clicks on a new color. This will cause the user to be re-centered at the new color by moving their wedge shape along with them. A connecting line is drawn as they move from color to color so that they can retrace their steps and visualize the path they are taking through the collection. Most commercial paint collections are not as dense in some locations making it difficult to move in a particular direction just by selecting colors. To alleviate this problem, the user is also able to click on the CIE LCH axes. This will re-center the focal point of the camera but not add a color to their collection or change the current color. It is simply a way of moving to a location where there is no predefined color.

The intended workflow within this module is as follows:

- The user scans their color.
- They examine the surrounding colors deciding for themselves which ones are closer to their desired color.
- The user adjusts the volume extent to see more or fewer colors in each direction.
- They decide on a color and double-click it to add it to their palette or they double-click on the axes to make small movements along the CIE LCH axis directions.
- Once they have centered on a color they can continue to make slight adjustments (tweaking it) by double-clicking other surrounding colors.
- This continues until they have found the one color that is closest to their intentions.

A second workflow was created around the idea of tweaking that works as follows:

- The user adds a color to their palette by some other means.
- They center the tweaker camera on this color by simply selecting it in their palette.
- The user can now proceed to tweak this color in the same manner as described above.

Since the colors are in CIE $L^*a^*b^*$ space, physical distance is a good approximation of perceptual distance. This enables the user to subconsciously approximate the Delta E. The ability to limit the range along the CIE LCH dimensions lends a certain intuition to the physical arrangement. For instance, if the customer would like a brighter or darker color they can create a very tall, narrow volume and see what colors lie along the lightness dimension. The system allows for the creation of volumes similar to this that isolate any of the dimensions or combine two dimensions at once.

The Color Tweaker was designed for looking at small, localized subsets of the full database of colors. Similarly, movements in this space are meant to be small and localized. It is possible, however, to zoom out and adjust the extent so that the entire collection is visible. Our system brings this visualization to the customer giving them a sense of the entire set of colors that was not present in the display rack or the fan deck.




Figure 6. The Color Navigator Module




Figure 7. The Color Navigator with the extent volume basket.

Color Navigation and Selection

There are two modules for the task of color navigation and selection, one for selecting colors with a standard appearance and another for selecting the metallic colors.

The Color Navigator

At present our color database is populated with all of the Benjamin Moore® colors (a set composed of over 3,300 different tints). The large size of this collection makes it difficult to filter through the database and locate a single color. In order to simplify this process, a module was created that is similar in spirit to The Color Tweaker module but provides a very different user experience and interface. *The Color Navigator module* (see Figures 6 & 7) is a *centralized* visualization of the color database again in CIE $L^*a^*b^*$ color space. It is intended for observing large subsets or the entire collection at once and selecting individual colors from this collection.

The Color Navigator does not draw individual colors but instead draws entire strips of colors. These strips are intended to correlate with the color chip cards that manufacturers make available on a store display rack or in a fan deck. The visualization is generated by fitting a 2nd degree polynomial to the $L^*a^*b^*$ values of the color cards in the fan deck. These cards can contain anywhere from 3 to 7 colors that usually represent tinting percentages away from titanium white (or a similar white base). This gives each card a somewhat constant hue and monotonic changes in lightness. The fitting of the cards to a parabolic curve ensures that they do not oscillate or curve excessively while still leaving them reasonably close to their measured location in the color space. The collection of strips is drawn above a circular platform with a map of the hue circle in-laid in its surface. This platform contains a set of controls that determines limits in each of the major perceptual dimensions: lightness, hue and chroma. By moving these limits the user can peel away colors on the outside or behind the area of interest revealing interior colors or isolating a section for closer examination.

As in The Color Tweaker, these limiting controls define an extent volume along the CIE LCH dimensions dictating how many strips and what portions of the strips are actually drawn. The density of the database and the relative closeness of colors (especially around the neutral axis) necessitates this ability to prune the size of the set. In addition, this filtering helps overcome the inherent obstruction of interior data points present in any three-dimensional visualization. When the user adjusts these

limiting controls a metallic wire frame is drawn indicating the outline of this volume extent (see Figure 7). It is similar in appearance to a basket and its function is similar in the sense that only the colors that are within the basket are drawn. If only part of the strip is within the volume then the colors inside it are drawn and the rest of the strip is drawn as a thin gray line. This will help the user understand where the rest of the colors on the strip lie relative to the color of interest. Strips that fall completely outside the basket are not drawn at all.

The camera's viewpoint is always centered on the CIE $L^*a^*b^*$ neutral axis. Unlike The Color Tweaker, the user cannot change the central position. As with all objects in our system the strands and platform can be moved all at once around the screen and re-positioned for easier viewing but the center of rotation for The Color Navigator is always the neutral axis. Initially the color strips are drawn as thin strands with a square cross-section so that the outer strips do not obscure the inner ones too much. When the mouse is placed over a particular color strip the strand is stretched so that it takes on the appearance of a card. In addition, an enlarged view of the colors that are present on the card is drawn over to the side as part of the surrounding user interface. This makes it easier to evaluate the color of interest. The user can zoom in and out to see the fine details of each color strip and they can select colors to add to their palette by double-clicking them.

The Studio Finishes®

A popular trend among household paint is to offer special latex based finishes that possess metallic and pearlescent reflective properties. One such product is the set of metallic colors in the Studio Finishes® collection from Benjamin Moore®. This is a relatively small collection of tints as are products offered by most household paint manufacturers that possess metallic or pearlescent qualities. Consequently, the set composed entirely of these colors will be much smaller than the full color database. Furthermore, it is impossible with standard color models and visual difference metrics to quantify simple perceptual relationships between these colors because they do not possess a single dominating tristimulus value. Because of this, *The Studio Finishes®* module uses a simple two dimensional grid layout (see Figure 8). This module allows all of the metallic and pearlescent colors to be viewed at once (presently using the aforementioned metallic and pearlescent tints from the Benjamin Moore® Studio Finishes® collection). It is also very important for the user to be




Figure 8. The Studio Finishes[®] Navigator Module

able to evaluate the complex reflective properties of these paints during the selection process. When these paints are placed on at objects (such as the square paints chips available in any paint store) much of their appeal is lost. To accommodate this, the system has the ability to place each color on a three dimensional object. At present each color is shown on a torus (a donut shaped ring) though the system has the ability to use any 3D object. The torus has the advantage that any given view of the object creates a multitude of different angles between the surface and the viewer.

Implementation

What follows is a brief description of some of the technical considerations that went into the actual implementation.

Hardware - Display Devices

One of the most vital components to the success of the Wall of Inspiration is the display device. For this task we have chosen a high-end 70 inch display device. These displays offer an unusually high native resolution and can be stacked on top of one another or locked together side by side to create an array that approximates a single display. We utilize four of these devices in a 2 by 2 arrangement. The cubes are designed to have a very small amount of space between them (in the range of 1-3 mm) and the image is projected right to the edge of the device. These displays use LCOS technology (Liquid Crystal On Silicon) which is essentially a reflected form of LCD technology. They also boast a very wide monitor gamut due in part to the separate LCOS chips for each color channel (similar to 3-chip DLP). Since these devices are modular they are easier to transport than an array of conventional projectors and a separate screen. In particular, it is relatively easy to adjust the geometry of each individual cube so that the image lines up appropriately. It will also be possible, once our research is completed, to separate the cubes and use them individually for other purposes extending their value beyond our experimental work. Color calibration techniques appropriate for an additive color reproduction device have been utilized to adjust and characterize the video wall. Careful measurements were made to determine the chromaticity coordinates of the video wall primaries and to set the white point for the display device. This information was used to form the matrix necessary to convert XYZ tristimulus values of the Benjamin Moore[®] colors into the RGB coordinates of the display device. Separate gamma correction measurements were made for each RGB primary and these adjustments were applied

to the RGB values determined using the transformation matrix.

Hardware - Display Workstation

Development of the software for the Wall of Inspiration has been made easier because of the hardware chosen to drive the display devices. We drive all four display cubes with a single computer. This computer uses the relatively new bus technology called PCI-Express. Previously, the bandwidth of a conventional PCI bus was too small to support the amount of data that needed to be exchanged between main-memory and the graphics card. This gave rise to the specialized graphics card bus known as the Advanced Graphics Port (AGP). This bus allowed dramatically higher bandwidth communication between main memory and the video card but did so with the limitation that only a single AGP could be included on a motherboard. Consumers still desired to run multiple-monitor desktops but they could only have one AGP graphics card. To alleviate this AGP graphics cards were enhanced to include duplicate hardware for the final D to A conversion of the video signal as well as duplicate frame buffers enabling a second video output to be included on a single card. This technology is called dual head.

As bus technology has progressed, the PCI bus that previously was ill suited for graphics tasks has improved. In its latest inception, PCI-Express, it meets or exceeds the speeds achieved by the AGP bus. Of particular interest is the ability for PCI-Express to include several slots on a single motherboard thereby enabling the installation of more than one video card (two being the current norm). Furthermore, by putting two dual head video cards in one system you can drive up to four display devices from a single computer. This is what enables us to run the Wall of Inspiration with a single workstation.

Being able to do this has many benefits. First, we avoid the need to transfer large amounts of data across a comparatively slow network connection. This keeps our rendering capabilities quite high. Second, we avoid the need to synchronize separate graphics card rendering clocks. Since the two cards are timed off the same bus their clocks will automatically be synchronized. Therefore, we can use consumer grade graphics cards as opposed to workstation grade and save cost. Third, we do not need to invest in large, complex software for coordinating the rendering process of several machines. The cost and overhead of maintaining this software is quite significant and eliminating it eases the budget considerably. Fourth, the software run on The Wall does not need to be specially designed to work in a distributed manner. In fact, The Wall can run existing, conventional 2D or 3D applications without modification. The only difficulty was creating a true fullscreen rendering environment which was achieved using standard Win32 API functionality.

Software - Design & Development

In a traditional power-wall system the software must be custom designed around the network of servers, the storage bays and the systems driving the projectors [1]. While a significant infrastructure of drivers and libraries are publicly available, they are designed around Unix-based platforms, are often not robust enough for commercial use, and require significant customization for the particular installation. Furthermore, the synchronization of several computers to a single rendering task requires significant amounts of communication and complexity just to divide the scene to be rendered into individual pieces. For this task there is commercial software available such as CAVELib. However, the added complexity of this software as well as its significant price tag make it an undesirable option. Since our system

is driven by a single computer directly attached to all of the display devices there is no need to coordinate a network of servers and storage or synchronize several projection systems. All of this overhead is inherently gone. Furthermore, we can write our rendering software using conventional OpenGL libraries without consideration for segmenting the scene. This is done for us in the system drivers for the graphics cards. Consequently, our software design has been a very simple, traditional, object-oriented approach. We augment OpenGL's capabilities by using Nvidia's Cg shading language in order to achieve sophisticated real-time lighting effects.

Software - Rendering Engine

The appearance of the paint color is computed throughout our system using a set of high dynamic range environment maps that have been pre-filtered using specific Phong cosine lobes as blur kernels similar to Heidrich and Seidel [2]. A filtered environment map was determined for the diffuse contribution that remains consistent across all paint colors. Other differently filtered environment maps exist for each of the five gloss levels: matte, eggshell, pearl, semi-gloss and high gloss. The parameters for the Phong reflection model are determined based on gloss measurements as described by Westlund and Meyer [4]. These environment maps are all pre-loaded to the video card so that the user can switch between them instantaneously. The appearance of the metallic and pearlescent paints is approximated by fitting the angular dependent spectral distribution of each color to a set of Phong cosine lobe basis functions each of which is again represented by a pre-filtered environment map. The process is similar to the one used by Shimizu and Meyer [3]. The basis functions are recombined using a Cg fragment shader program on consumer level graphics cards thereby enabling the entire system to run in real-time.

Spectral curves were available for all paint colors in our system. This data was used to compute XYZ tristimulus values in the conventional manner. This data is then converted to other color spaces (such as CIE L*a*b*, LCH and various device dependent RGB color spaces). All conversions are done using standard colorimetric procedures based on physically sampled values whenever possible. Tristimulus data was generated for a variety of standard light sources (D65, A, F2) as well as some proprietary light sources developed by Benjamin Moore that approximate a combination of daylight and tungsten illuminates (a common situation in a home). All of these values were precomputed and made available in a database.

Conclusions

New methods for selecting and evaluating the color of paint were created by taking advantage of emerging developments in computer graphics software and hardware. A hue, saturation, and brightness interface was provided for navigating through the traditional color cards that define the manufacturer's paint collection. The position of a color sample, brought to the store by the consumer, could be visualized within the collection and either visually matched or subtly adjusted to meet the customer's satisfaction. Solid colors could be displayed, including optional gloss levels, on complex surfaces. One of the major improvements of the system over existing color selection tools is the ability to evaluate the reflection properties of glossy paints and newer effect finishes that have metallic and pearlescent properties. The end result is a software package consisting of individual modules that work together seamlessly. The pictures generated by the software are displayed on a large four panel video wall. The

increase in size potentially makes it possible for the customer to better evaluate the appearance of the color when it is painted on an actual wall.

Acknowledgments

The authors would like to acknowledge Pete Wagner for his participation during the initial conceptual discussions of this project. We would also like to thank Benjamin Moore for support and funding of this project.

References

- [1] Alex Elder, Thomas M. Ruwart, Benjamin D. Allen, Angela Bartow, Sarah E. Anderson, and David H. Porter. The intensity powerwall: A case study for a shared file system testing framework. *IEEE/NASA Goddard Symposium on Mass Storage Systems*, pages 151–161, 2000.
- [2] Wolfgang Heidrich and Hans-Peter Seidel. Realistic, hardware-accelerated shading and lighting. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 171–178, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.
- [3] Clement Shimizu and Gary W. Meyer. Computer aided color appearance design using environment map based lighting. *Computational Aesthetics in Graphics, Visualization and Imaging*, pages 223–230, 2005.
- [4] Harold B. Westlund and Gary W. Meyer. Applying appearance standards to light reflection models. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 501–511, New York, NY, USA, 2001. ACM.

Author Biography

Seth Berrier received his BA in Music from Butler University in Indianapolis. He received his MS in Computer Science from the University of Minnesota where he is currently pursuing a PhD in the same. His research interests include photographic acquisition of reflective properties, computer aided art design, numerical color harmony and music visualization.

Benjamin Moore, Personal Color Viewer and Studio Finishes are registered trademarks licensed to Benjamin Moore & Co.