Superimposing Pictorial Artwork with Projected Imagery

Oliver Bimber, Franz Coriand, Alexander Kleppe, Erich Bruns, Stefanie Zollmann, and Tobias Langlotz
Bauhaus University Weimar

We present a novel approach for using pictorial artwork as information displays and show how to combine almost any kind of computer-generated visual information directly with the painted content.

Working high above the floor of the Sistine Chapel in the Vatican of Rome, between 1509 and 1512 Michelangelo Buonarroti painted some of the finest pictorial images of all time. On the ceiling of the papal chapel, he created a masterpiece fresco that includes nine scenes from the book of Genesis. Among them is the famous *Creation of Adam* scene—showing God touching Adam's hand. In 1510, an initial study led Michelangelo to draw the Adam figure as a sanguine on a piece of paper. Today, this early drawing is displayed at the London British Museum.

Around 1518 Jacopo Pontormo painted yet another scene from the book of Genesis, called *Joseph and Jacob in Egypt*. It decorated the bedchamber of Pier Francesco Borgherini for many years and is now being displayed at the London National Gallery. This oil painting made headlines after art historians discovered incredible underdrawings beneath the top paint layer. The underdrawings show that the artist had decided to flip the background design after starting work and simply overpainted his initial approach.

Such underdrawings have been found in many Renaissance and Baroque paintings. In 1634, for example, Rembrandt van Rijn painted a self-portrait that was later retouched by one of his students to feature a Russian nobleman. The original portrait was hidden under layers of paint for more than 300 years, until it was uncovered recently by art restorers. It sold for nearly seven million British pounds.

Pictorial artwork, such as these examples, can tell interesting stories. The capabilities of museums to communicate this and other information, however, are clearly limited. Text legends and audio guides can mediate facts, but offer little potential for presenting visual content such as embedded illustrations, pictures, animations, and interactive elements.

Also because of the difficult economic situation, edutainment is becoming an important factor for museums. By applying new media technologies—such as computer graphics, virtual reality, and augmented reality—exhibit-oriented information might be communicated more effectively, but certainly in a more exciting way.

In this article we describe a novel technological approach, a mathematical model, a real-time rendering algorithm, and examples of presentation techniques for integrating almost any kind of visual information directly into pictorial artwork. Our system displays such information while keeping the observers' attention on the original artifact and doesn't require additional screens.

Technical approach

A seamless and space-efficient way for integrating visual information directly into pictorial artwork is to use the artwork as an information display (see Figure 1). The display can serve as a diffuse projection screen, and we can apply conventional video projectors to show computer graphics together with the painted content. To perceive the projected imagery in the correct colors and intensities, however, requires that the influence of the underlying physical color pigments is neutralized. In most situations, this isn't possible if untreated images are simply projected directly onto arbitrarily colored surfaces. The problem is that the projected light interacts with the color pigments on the canvas and is partially absorbed if the pigment's color isn't fully white.

To solve this problem, we use a new film material that has two properties: first, it's completely transparent and second, it diffuses a fraction of the light projected onto it. The film consists of an even deposition of fine particles on both sides of a polyester base with no visible artifacts. It was used for creating special effects in Hollywood movies such as *Minority Report* and *Paycheck* and sells for $350 per square foot. Initial

1070-986X/05/$20.00 © 2005 IEEE

Published by the IEEE Computer Society

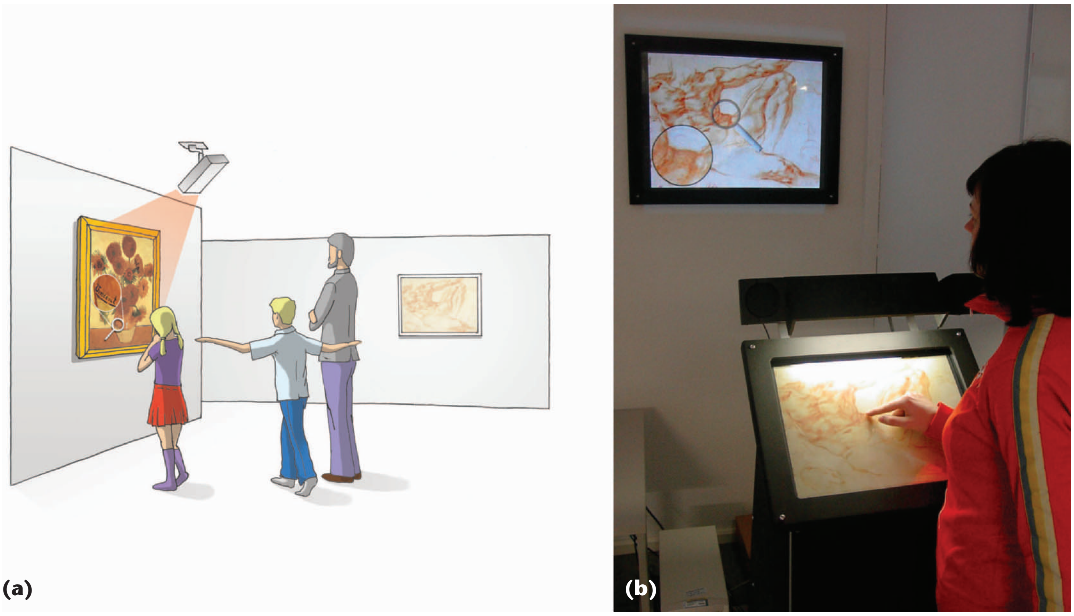

Figure 1. Using pictorial artwork as information displays: (a) concept sketch and (b) first prototype.

measurements have revealed that on average 20 percent (± 1 percent) of the light that strikes the film is diffused while the remaining fraction is transmitted toward the canvas. This 0.1-mm thin, transparent film can be seamlessly overlaid on the canvas (with or without direct contact) by integrating it into the frame that holds the artwork. We use BenQ 7765PA 1,100 ANSI lumens extended graphics array (XGA) digital light projectors to display images on film and canvas.

Mathematical model

In this section we describe how light interacts with a textured surface through the film. If a light beam with incident radiance L is projected onto the transparent film material located on top of the original artwork, a portion d of L is directly diffused from the film while the remaining portion t of L is transmitted through the film. The transmitted light tL interacts with the underlying pigment's diffuse reflectance M on the canvas, and a color-blended light fraction tLM is diffused. The portion $tLMt$ is then transmitted through the film, while the remaining part $tLMd$ is reflected back toward the canvas, where it is color blended and diffused from the same pigment again. This ping-pong effect between film material and canvas is repeated infinitely while for every pass a continuously decreasing amount of light is transmitted through the film that contributes to the resulting radiance R. Mathematically, we can express this as an infinite geometric series that converges toward a finite value. The same is true for the environment light with incident radiance E emitted from uncontrollable light sources. Because these light sources also illuminate the canvas and the film material, we must also consider the environment light's contribution to R.

Figure 2a (next page) shows this process as a sequence diagram. Note that in contrast to this conceptual illustration, normally no physical gap exists between the film material and canvas, and the light interaction occurs at the same spot.

If all parameters (L, E, M, t, and d) are known, we can compute the resulting radiance R that's visible to an observer in front of the canvas:

$$\begin{aligned} R &= (L+E)d + (L+E)t^2M\sum_{i=0}^{\infty}(Md)^i \\ &= (L+E)\left(d + \frac{t^2M}{1-Md}\right) \end{aligned} \tag{1}$$

Now that we know R,which we expect to see, we need to solve Equation 1 for L:

$$L = \frac{R}{\left(d + \frac{t^2M}{1-Md}\right)} - E \tag{2}$$

Equation 2 allows computing the incident radiance L that needs to be projected onto the film and the canvas to create the known result R. The radiant intensity I of the projector to create L is related to a discretized pixel value and is given by

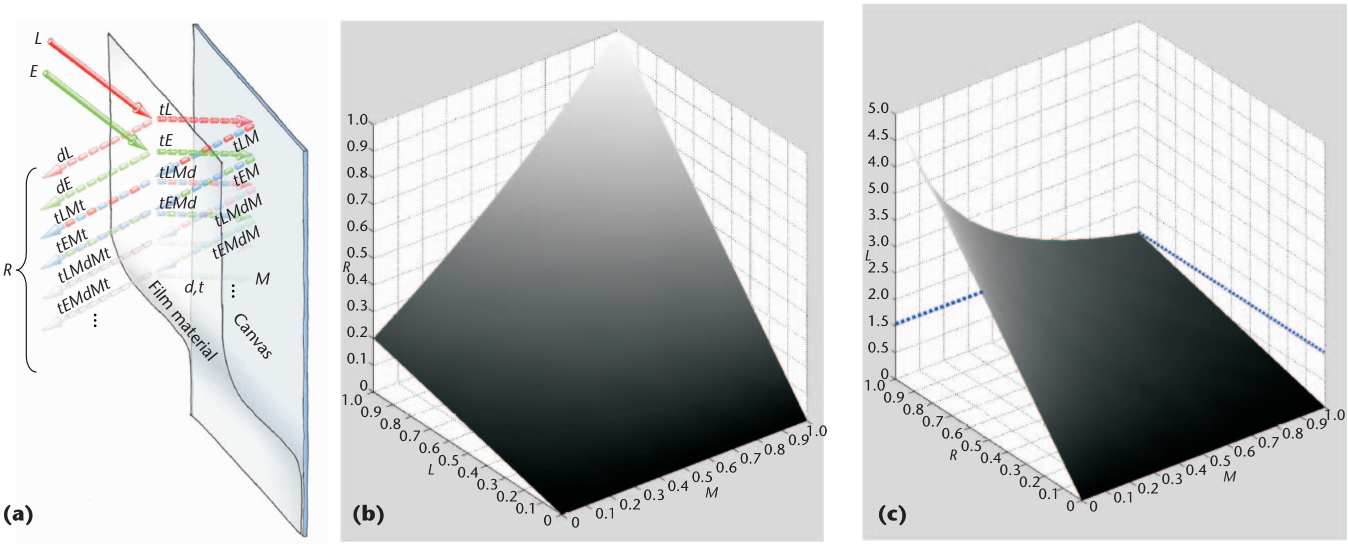

Figure 2. Properties of the film material. (a) Interaction of projected light and environment light with the canvas and the film material (sequence diagram). Visualization of (b) Equation 1 and (c) Equation 2 over R, L, and M (without E, t = 80 percent and d = 20 percent).

$$I = L\frac{r^2}{\cos\alpha}s \qquad (3)$$

where $r^2/\cos\alpha$ are the form factor components: square distance attenuation and angular correlation of the projected light onto the canvas. The additional factor s allows scaling the intensity to avoid clipping and to consider the simultaneous contributions of multiple projectors.

Our approach has clear limitations, which are illustrated in Figures 2b and 2c. Not all radiances R can be produced under every condition. If M is dark, most of L and E are absorbed. In an extreme case, the corresponding pigment is black ($M = 0$). In this case the right term of Equation 1 is canceled out. The remaining left term—which depends on the diffusion factor d of the film material—sets the boundaries of the final result that can be produced. The intersection of the surface with the RL-plane in Figure 2b illustrates these limitations. Consequently, in the worst case of our example, only 20 percent of R can be generated. This situation is also reflected in Figure 2c as the intersection of the surface with the LR-plane. Here we want to assume that $sr^2/\cos\alpha = 1$, which results in $L = I$. For a single video projector, the projected radiance L and consequently the radiant intensity I cannot exceed the normalized intensity value of 1 (dotted line). But for creating most of the resulting radiance values, L and I must be larger. This situation worsens for $r^2/\cos\alpha > 1$ and for $E \to 1$ or $M \to 0$.

However, the contributions of multiple (n) projectors allow displacing this boundary with

$$L = \sum_{i=1}^{n} L_i, \quad I_i = L_i\frac{r^2}{\cos\alpha_i}s_i \qquad (4)$$

If all projectors provide linear transfer functions (for example, after gamma correction) and identical brightness, $s_i = 1/n$ balances the load among them equally. However, s_i might be decreased further to avoid clipping and to consider differently aged bulbs.

For Figures 2b and 2c, we don't consider the environment light E and set it to zero. Additional environment light would simply shift the surface in Figure 2b up on the R-axis, and the surface in Figure 2c down on the L-axis. Note that our mathematical model must be applied to all color channels (such as red, green, and blue for projected graphical images) separately.

Real-time color correction

We implemented Equations 2 through 4 as a pixel shader to support real-time color correction. We used Nvidia's Cg framework for fragment processing on an MSI GeForce FX 5600 TD graphics board. Figure 3 illustrates the rendering process based on an example of Michelangelo's *Creation of Adam*. Although we'll use this example to explain the process, it's universal and can be applied with arbitrary background images.

In our example, a copy of the original Adam drawing serves as a background image. Our goal is to overlay it entirely with a registered view on the actual ceiling fresco of the Sistine Chapel.

The first step of the rendering process is to create an input image I_i, which needs to be overlaid.

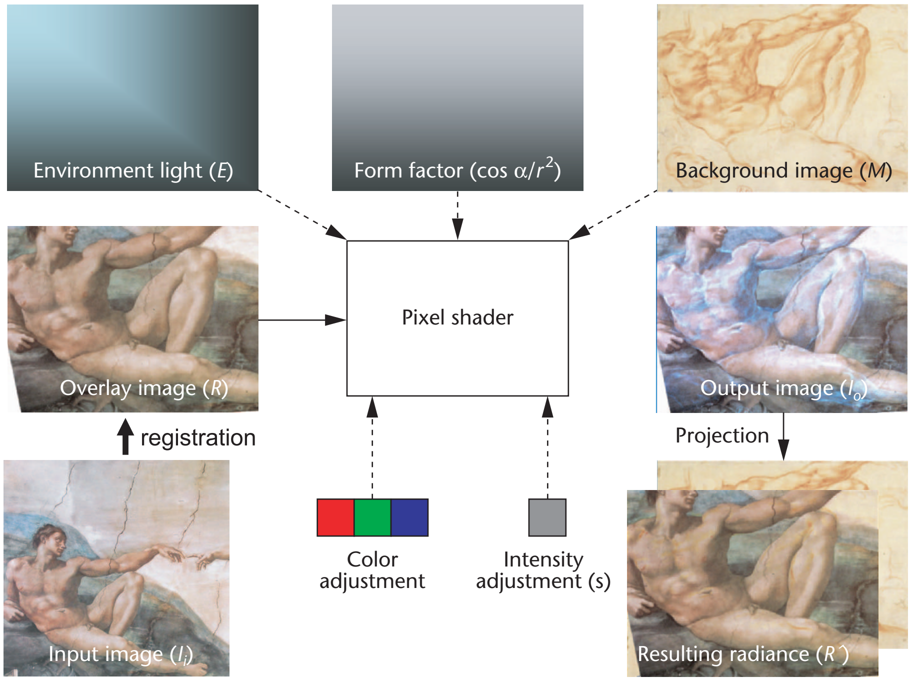

Figure 3. Real-time color-correction process.

This image can be dynamically rendered (as part of a real-time animation or an interactive experience), played back (frames of a prerecorded movie), or static (a photograph of the corresponding ceiling portion—as is the case in our example).

The input image must be registered to the physical drawing. Registration is achieved by texture mapping I_i onto a predistorted image geometry that's precisely aligned with the physical drawing. The amount of distortion depends on the geometric relation between the video projector and the canvas. The distortion can be simple keystone deformations to more complex curvilinear warping effects (for example, if lens distortion of the projector has to be neutralized).

Several automatic approaches apply video cameras and image analysis to align multiple images of tiled projection screens.[1,2] A structured light registration benefits from controllable features, such as projected grid edges or Gaussian matchpoints that can be easily detected in the camera views with a subpixel precision.

In our case, however, we have to register a digital representation of the artistic content against its physical representation on the canvas, rather than registering one projected structured light image against another. To detect nonstructured artistic features (such as fine lines) in the artwork and register them automatically against the corresponding features in the digital content represents an important task of computer vision—especially if projected pixels and physical pigments must align precisely on the canvas. We're critical about the feasibility and precision of an automatic method for our problem and have decided to solve it with a manual registration process that benefits from the resolution of the human visual system. Because the following steps have to be performed only once, we believe that they represent an acceptable solution.

Our manual registration process lets users interactively identify 2D correspondences between artistic features in the background image M within the image space—that is, an image of the drawing displayed on a control screen—and the physical drawing on the wall. This is done using a 2D input device, such as a conventional mouse whose pointer is visible on the control screen and as projection on the canvas.

We use dual output graphics card and an additional video signal splitter to drive the control screen and one or two projectors. The result of this feature matching is a set of 2D vertex fiducials with their corresponding texture coordinates within the image space. The fiducials are Delauny triangulated, and the system uses texture coordinates to map the correct image portions of I onto the image geometry. This results in the overlay image R. Here we should stress again that a precise correspondence between R and M is important to achieve qualitatively good results.

In our experiments, the measurement of 50 to

70 fiducials proved sufficient. In contrast to uniform grid methods normally applied for projector alignment, this general geometric registration allows correlating arbitrary artistic features in the physical drawing with the corresponding pixels of M in the image space.

Thus, it provides an accurate matching that can be regionally improved further if linear interpolation within single-grid triangles fails to be precise enough. The registration results don't change if the projector and background image are fixed. Before we can render R, we must enable the color-correction pixel-shader. Five parameters are passed to the pixel shader to ensure that Equations 2 through 4 can be computed.

The first parameter is the environment light E in the form of intensity texture that has the same size as R. It contains intensity values that represent the uncontrollable lighting situation on the canvas. We can determine the intensity values by measuring the irradiance of the environment light with a light meter for a discrete number of sample spots on the canvas' surface—resulting in the lux (lx) values E'. These values have to be normalized to an intensity space that ranges from 0 to 1 so the shader can process them. To do this, we measure the same spots again, but this time the highest intensity possible (for example, a white image) is projected onto the light meter, which is measured in addition to the environment light. These measurements are equivalent to the total irradiance $T' = L' + E'$, and also carry the unit lux.

Because we know that $L' = T' - E'$ is equivalent to the scaled intensity value $\cos \alpha/r^2$, we can convert the measured radiance of the environment light from lux into the normalized intensity space with $E = E'/(T' - E') \cos \alpha/r^2$. To approximate the intensity values for the entire image area in E, all the measured spots are mapped onto the image space and are Delauny triangulated. The values for the remaining pixels are linearly interpolated by the graphics card while rendering the Delauny mesh. Note that E is constant if the environment light doesn't change. For reasons that we describe next, we can assume that $\cos \alpha/r^2$ is constant and equals 1.

The second parameter is the form factor that represents the geometric relation between the video projector as a point light source and the canvas. Because it doesn't change for a fixed relation between projector and canvas, we can precompute it and pass it to the pixel shader in the form of an intensity texture with the same dimensions as E and R.

Like Raskar's Shader Lamps,[3] the graphics pipeline can produce this texture—a geometric model of the canvas is rendered with a white diffuse reflectance from the viewpoint of the projector. Attaching a virtual-point light source (also with a white diffuse light component) to the position of the projector and enabling square distance attenuation produces intensities proportional to $\cos \alpha/r^2$. The pixel shader can compute the required reciprocal. Practically (that is, for normal-sized canvases and nonextreme projector configurations), we found that the form factor is constant over all pixels. It's then contained by the intensity adjustment parameter s.

The third parameter is the background image M. It also has the same dimensions as E and R. This image can be generated by, for example, scanning the color values or taking a photograph of the original drawing under uniform illumination.

The fourth and fifth parameters contain color and intensity adjustment values that let users fine-tune the video projector's individual color response and prevent intensity clipping. These parameters also help users adopt color drifts that they can introduce while capturing the background image and allow users to consider the contributions of multiple projectors. Note that gamma correction must be applied in advance. This mostly occurs with projectors with nonlinear transfer functions as well as projectors with linear transfer functions that apply a de-gamma mapping on the video input. Gamma correction is usually supported by the graphics hardware and the video driver, but the pixel shader can also carry it out. We adjust these values manually, but the support of automated methods for color matching multiple projectors[4] is imaginable.

The output image I_o is the final result of this rendering process and will be displayed by the video projector. If projected geometrically correctly onto the drawing, the result R' will be visible. Both images—R and R'—are mostly identical, except for slight artifacts that appear because of previously discussed limitations. Figure 4 shows the results of our example projected onto a real drawing with a single video projector.

Apparently, the underlying drawing can be made partially or completely invisible to display the graphical overlay in the correct colors on top of it. Figures 4e through 4h show close ups in which diverging body parts (such as belly and knee) are overdrawn and displaced by the projection.

Some intensities and colors that are required

Image courtesy of the British Museum, London

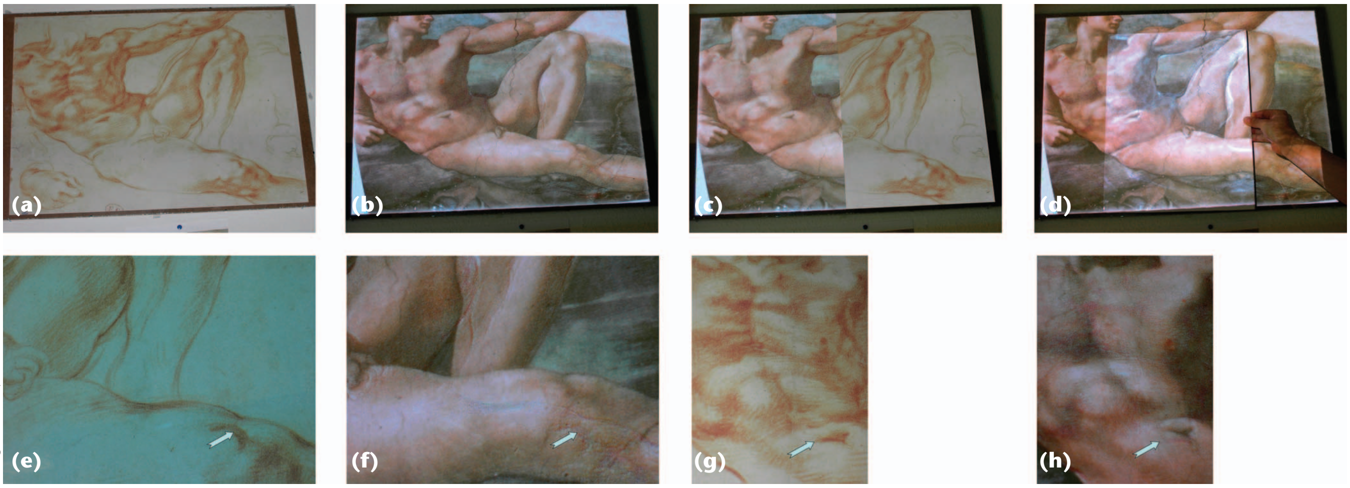

Figure 4. Results of the color-correction process with a single projector on a real drawing. (a) Real drawing (64 × 48 cm) under environment light. (b) Output image emitted onto drawing. (c) Partially augmented drawing. (d) Output image on a white piece of paper. (e–h) Close ups. While the upper body coincides with the drawing and painting, Michelangelo modified the lower body. The arrows indicate the displaced knee and belly sections. They point at the same spot on the drawing.

to neutralize the underlying color pigments can't be achieved by a single video projector. The worst case is to turn a black pigment on the canvas into a white color spot. Figures 2b and 2c illustrate that in such a case the required intensity can easily exceed the boundary of 1 in our normalized intensity space. The pixel shader clips these values to 1, which results in visible artifacts.

The simultaneous contributions of multiple projectors can reduce or even eliminate these effects. Figure 5 shows the extreme case of an input image that has no geometric correspondences to the underlying background image. In addition, both projections together create bright colors (the sky) on top of dark color pigments on the canvas. In Figure 5a, a single projector is used. Intensity values that are too large are clipped and result in visible artifacts. Balancing the intensity load between two projectors reduces these artifacts clearly (see Figure 5b). Figures 6 through 8 (next page) show more results for other cases.

Because of the hardware acceleration of today's graphics cards, we can easily perform the color-correction process in real time. Note that none of the photographs in this article have been retouched. Slight variations in color and brightness are because of different camera responses.

Image courtesy of the Vatican Museum, Rome

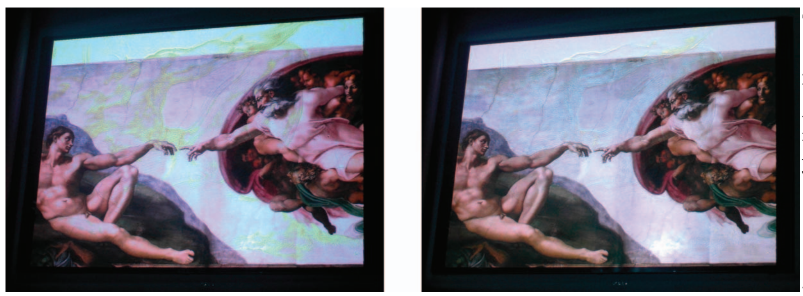

Figure 5. Results of color-correction process with two projectors. (a) Limited intensity capabilities of a single projector result in visible artifacts. (b) The contribution of a second projector reduces these effects.

Content creation, authoring, and presentation

The basic color-correction process allows visualizing all sorts of information and effects by rendering the desired result as input image I_i into the pixel shader. This opens the potential for applying a wide range of established or new presentation tools and techniques. We want to categorize them into six main classes:

- inlay objects (textual information, images, videos, and arrow annotations);
- lens effects (magnification, x-ray, toolglasses, and magic lenses[5]);
- focus effects (blurring, decolorization, hiding, and highlighting);
- 3D effects (walk- or flythrough and object observation);
- modification effects (drawing style modifica-

Image courtesy of the Vatican Museum, Rome

Figure 6. Sample scenes of the papal chapel's ceiling projected onto the Adam drawing (Figure 4a) during an interactive slide presentation. The underlying drawing is barely visible.

Image courtesy of the National Gallery, London

Figure 7. Pontormo's* Joseph and Jacob in Egypt *(65 × 49 cm). (a) Copy of original painting illuminated under environment light, (b) modification of painting style from oil on wood to watercolor on paper via a 2D artistic filter, (c) re-illumination and lens flare, and (d) registered visualization of underdrawings (infrared recordings are black and white).

tion, re-illumination, and color restoration[6]); and

- audio effects (verbal narrations, music, and sound that are synchronized to the visual effects).

Figure 9 illustrates a variety of examples. Embedded information windows (see Figure 9a), such as the semitransparent text panel or the opaque image and video panel are—in contrast to simple physical text labels—dynamic and can be directly linked to corresponding portions of the artwork (via arrow annotations, for example). The magnifying glass in Figure 9b allows zooming into interesting areas to identify brush strokes and hatching styles of the artist. Usually museum visitors are restricted from getting close enough to recognize such details. To draw the observers' attention to a certain image portion, it can be brought into the foreground while the remaining part is eclipsed. In Figure 9c, the focus area is highlighted while the rest is decolorized.

This is technically achieved by projecting the inverse colors of the background image onto *M*. In this case, the colors of the pigments on the canvas are physically canceled out and *M* appears in grayscale. Figure 9d shows a 3D flythrough in the Sistine chapel—building a spatial correlation to surrounding paintings and the environment. This allows relocating the observer virtually to remote places.

Stereoscopic rendering (and optional head-tracking technology) allows an immersive experience. Although the transparent film preserves the polarization of light, the underlying canvas doesn't. Consequently, we can only apply active or anaglyphic stereo solutions. We use a mechanical (half opaque and half transparent) shutter wheel rotating with 120 Hz in front of two video projectors. A light sensor measures the wheel's position and triggers the shutter glasses' infrared synchronization signal. In combination with conventional LCD projectors, this is a simple and—compared to CRT projectors—cost-efficient alternative for active stereo.

Figure 7 illustrates further presentation examples, such as modifications of the painting style via 2D artistic filters, visualization of the underdrawings, and scene modification through re-illumination.

For noninteractive presentations, it's possible to pregenerate the entire content. In this case, we can apply well-established content creation and authoring tools (such as digital imaging, 3D modeling, video editing, and audio-mixing packages). These tools already provide techniques such as animations, image filtering, rollover effects, and so on, as well as professional graphical user interfaces. For interactive presentations, however, the generated content must be managed by the presentation software linked to an interaction framework.

If no user interaction is required, we apply an embedded Audio Video Interleaved (AVI) player to map video frames onto input images I_i on the fly. This represents a direct interface to our own player framework that comprises rendering and color correction, and allows the user or content creator to benefit from the capabilities of established content creation and authoring packages, such as Adobe Photoshop, Discreet 3ds max, Alias Maya, or Adobe Premiere. Our content interface for 3D interactive and stereoscopic/head-tracked presentations is Nvidia's NVB format, which can be exported from Discreet's 3ds max and displayed by our player.

To see how others have approached incorporating projectors, see the sidebar on the next page, "Previous and Related Work."

Future work

Using pictorial artwork as information displays opens another door to embedded multimedia content in museums and art galleries. Our proposed method is simple, seamless, cost and space efficient, robust, and compatible with off-the-shelf hardware and software. These are all important factors for museum operators, content creators, and museum visitors. The presented techniques let us think of a wide variety of presentation and interaction tools that we didn't explore in this article. Dynamic view management and label placement are only two examples. As for other areas, the development of efficient interaction techniques and devices for such displays will be an interesting challenge.

As we previously discussed, our method has limitations that are mainly defined by the capabilities of the applied hardware. For example, the restricted resolution, brightness, contrast, minimum focus, distance, and black level of video projectors are issues that will certainly be improved by next-generation projectors. The XGA resolution and the brightness of 1,100 ANSI lumens of our low-cost projectors were appropriate for small- and medium-sized paintings in a normally lit environment. An upscaling is possible by using more projectors, but downscaling would either result in a loss of effective resolution or in focus problems.

Black, for instance, is a color that can't be projected. Instead the environment light together with the black level of the projectors illuminates areas that need to appear black. We found that

Image courtesy of the Museum het Rembrandthuis, Amsterdam

Figure 8. Rembrandt's self-portrait (48 × 56 cm). (a) Copy of original painting as it looks today (illuminated under environment light). Various cleaning stages to remove the overpainted layers from (b) 1935, (c) 1950, and (d) 1980 are projected onto (a). Only black-and-white photographs of these stages are available. The high black level of the video projectors prevents us from creating a totally black color on the canvas. Extreme regions, such as overlaid hair and a hat can't appear completely black for this reason.

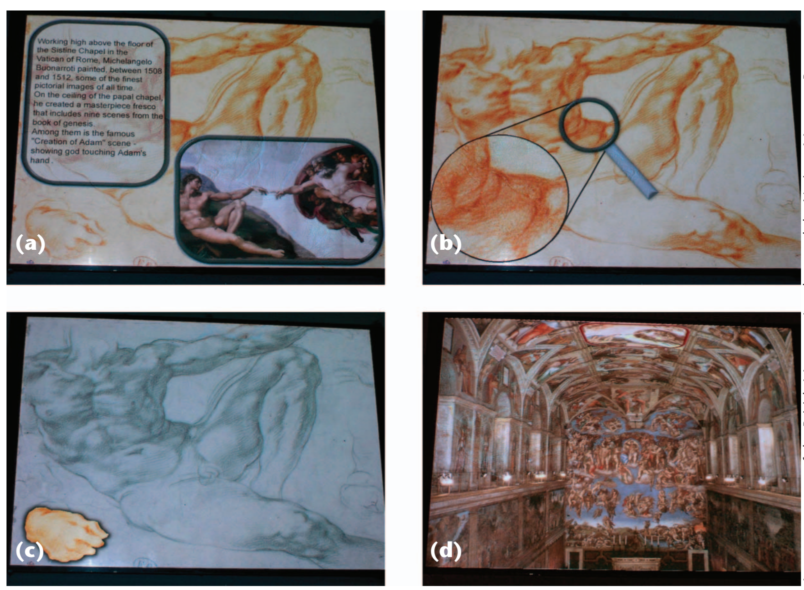

Images courtesy of the British Museum, London, and the Vatican Museum, Rome

Figure 9. Examples of presentation techniques: (a) inlay text and image, (b) magnification, (c) focus through highlighting and decolorization, and (d) 3D flythrough in the Sistine chapel. Note that in images (a–c) the Adam drawing itself isn't projected.

Previous and Related Work

With increasing capabilities and decreasing costs, video projectors have become widespread and established presentation tools. Being able to generate images that are larger than the actual display device virtually anywhere is an interesting feature for many applications that can't be provided by desktop screens. Several research groups have discovered this potential by applying projectors in unconventional ways to develop new and innovative information displays that go beyond simple screen presentations.

The Luminous Room,[1] for instance, describes an early concept for providing graphical display and interaction at each interior architecture space's surface. Co-located two-way optical transducers—called *I/O bulbs*—that consist of projector-camera pairs capture the user interactions and display the corresponding output. With the Everywhere Displays Projector,[2] Pinhanez has extended this concept technically by allowing a steerable projection using a pan/tilt mirror. Later, Raskar et al.[3] demonstrate how context-aware handheld projectors—so-called *iLamps*—can be used as mobile information displays and interaction devices.

Raskar et al. also use multiple projectors for their Shader Lamps[4] approach to lift the visual properties of neutral diffuse objects that serve as a projection screen. The computed radiance at a point of a nontrivial physical surface is mimicked by changing the bidirectional reflectance distribution function (BRDF) and illuminating the point appropriately with projector pixels. Animating the projected images lets us create the perception of motion without physical displacement of the real object.[5] This type of spatial augmentation is also possible for large, human-sized environments, as demonstrated by Low et al.[6]

Projector-based illumination has become an effective technique in augmented reality to achieve consistent occlusion[7,8] and illumination[9] effects between real artifacts and optically overlaid graphics. Video projectors instead of simple light bulbs are used to illuminate physical objects with arbitrary diffuse reflectance. The per-pixel illumination is controllable and can be synchronized with the rendering of the graphical overlays. It also makes the combination of high-quality optical holograms with interactive graphical elements possible.[10] Using a video projector to produce a controlled reference wave lets us reconstruct a hologram's object wave partially—but not at those portions that are overlaid by integrated graphical elements.

Yoshida et al.[11] describe a virtual retouching technique that applies video projectors for reconstructing and enhancing the faded colors of paintings by projecting light directly onto the canvas. An affine correlation between projection and the result captured by a camera is established manually for each pixel. Users can then retouch the original painting interactively via a desktop GUI.

In the context of our own approach, we can divide these methods into four main groups:

- information displays that project arbitrary images onto surfaces with arbitrary reflectance, but don't consider the effect of color blending;[1-3]
- techniques that project colored images onto surfaces with a neutral white reflectance to avoid color blending;[4-6]
- methods that project a uniformly colored illumination onto surfaces with arbitrary reflectance and texture; [7-10] and
- approaches that project colored images onto surfaces under consideration of their reflectance.[11]

As with Yoshida's work, our approach belongs to the fourth group. Although there are basic conceptual similarities between both attempts, they differ in their general technological realization, aimed applications, as well as in their mathematical model and rendering techniques. We want to know whether a direct projection of colored light onto surfaces with arbitrary reflectance will work effectively under real-world conditions. The limitations of such an approach will be reached quickly in cases where an adequately large portion of the color spectrum is absorbed by the underlying surface and can consequently not be reflected.

References

1. J. Underkoffler, B. Ullmer, and H. Ishii, "Emancipated Pixels: Real-World Graphics in the Luminous Room," *Proc. ACM Siggraph*, 1999, ACM Press, pp. 385-392.

even in our case the human vision system adjusts well to local contrast effects—which makes these areas appear much darker than they actually are. Even with little environment light, however, the high black level of video projectors causes this illusion to fail in extreme situations, such as the one shown in Figure 9. The development of video projectors indicates that we can expect a decrease of the black level and an increase of the contrast ratio in the future.

Light can damage the artwork. Ultraviolet (UV) and infrared (IR) radiation produced by the lamps of video projectors is critical. Commercially available UV/IR blocking filters can be mounted in front of the projectors' lenses to remove most of these unwanted rays while transmitting visible wavelengths.

For the remaining visible light portion, a general rule of thumb in the museum field advises to illuminate valuable and delicate pieces perma-

2. C. Pinhanez, "The Everywhere Displays Projector: A Device to Create Ubiquitous Graphical Interfaces," *Proc. Ubiquitous Computing*, Springer, 2001, pp. 315-331.
3. R. Raskar et al., "iLamps: Geometrically Aware and Self-Configuring Projectors," *Proc. ACM Siggraph*, ACM Press, 2003, pp. 809-818.
4. R. Raskar et al., "Shader Lamps: Animating Real Objects with Image-Based Illumination," *Proc. Eurographics Rendering Workshop,* Eurographics Assoc., 2001, pp. 89-102.
5. R. Raskar, R. Ziegler, and T. Willwacher, "Cartoon Dioramas in Motion," *Proc. Int'l Symp. Nonphotorealistic Animation and Rendering*, ACM Press, 2002, pp. 7-ff.
6. K. Low et al., "Life-Sized Projector-Based Dioramas," *Proc. Symp. Virtual Reality Software and Technology*, IEEE CS Press, 2001, pp. 93-101.
7. O. Bimber and B. Fröhlich, "Occlusion Shadows: Using Projected Light to Generate Realistic Occlusion Effects for View-Dependent Optical See-Through Displays," *Proc. IEEE/ACM Int'l Symp. Mixed and Augmented Reality*, IEEE CS Press, 2002, pp. 186-195.
8. S. Noda et al., "An Optical See-Through Mixed Reality Display with Real-Time Rangefinder and an Active Pattern Light Source," *Trans. Virtual Reality Soc. of Japan*, Virtual Reality Soc. of Japan, vol. 4, no. 4, 1999, pp. 665-670.
9. O. Bimber et al., "Consistent Illumination within Optical See-Through Augmented Environments," *Proc. IEEE/ACM Int'l Symp. Mixed and Augmented Reality*, IEEE CS Press, 2003, pp. 198-207.
10. O. Bimber, "Combining Optical Holograms with Interactive Computer Graphics," *Computer*, vol. 37, no. 1, Jan. 2004, pp. 85-91.
11. T. Yoshida, C. Horii, and K. Sato, "A Virtual Color Reconstruction System for Real Heritage with Light Projection," *Proc. Virtual Systems and Multimedia*, IEEE CS Press, 2003, pp. 158-164.

nently with no more than 100 lx–150 lx. The potential damage caused by light is cumulative (for example, 1 hour with 1,000 lx equals 1,000 hour with 1 lx). The damage also depends on the material and color of the painting and the wavelength (color) of the light. A temporary illumination of a higher light intensity isn't critical.

To give rough indications of the light levels that were reached in our experiments, we measured an ~300 lx maximum for the results shown in Figure 4 (single projector), an ~400 lx maximum for the results shown in Figure 6 (two projectors), and an ~700 lx maximum for the results shown in Figure 5 (also two projectors).

During a 2–3 minute presentation, this increased lighting is only temporary and such highlight situations usually only appear (if at all) for a short period of time (a few seconds) at varying locations on the canvas. Nevertheless, using the intensity adjustment we've described, the maximum light level can be constrained to be below an upper threshold. However, this might cause visible artifacts depending on the presented content, the painting, and the environment light (as described in Figure 2). Thus, it's important to reach a good balance between total illumination (projected and environment light) over time and creating convincing presentation effects.

We currently consider only the intensity of the environment light *E*. This is adequate for regular white light sources but will result in artifacts if visible color shading is created on the canvas by the environment illumination. Without modifying the mathematical model or rendering process, the environment light's color can be compensated by determining it with the aid of colorimeters, encoding this information in *E*, and passing it to the pixel shader.

Using cameras in combination with projected structured light samples might help us automate our registration, color, and intensity adjustment steps. This, however, requires that an acceptable level of precision can be achieved.

The presented concept and techniques are applicable in combination with diffuse pictorial artwork, such as watercolor paintings, pen or ink drawings, sanguine sketches, or matte oil paintings. Extreme light and view-dependent effects, such as non-Lambertian specular reflections; self-shadows; subsurface scattering and interreflections that are created by brush strokes, paint material; or canvas textures can't be handled yet. We'll investigate these ideas in our future research. However, some of these cases will require components—such as head tracking—that might not be effectively integrated into museum environments.

We plan to evaluate our approach in museums and art galleries. We'll have to conduct user studies and surveys on the efficiency and acceptability of this technology in the future. The feedback from visitors, curators, and staff members will lead to improvements and new ideas for the interactive presentation of pictorial artwork. **MM**

Acknowledgments

We thank the Museum het Rembrandthuis, Amsterdam, for providing digital photographs used for the experiment shown in Figure 8, and Matthias Hanzlik for the illustrations in Figures 1 and 2.

References

1. Y. Chen et al., "Automatic Alignment of High Resolution Multi-Projector Displays Using an Uncalibrated Camera," *Proc. IEEE Visualization*, IEEE CS Press, 2000, pp. 125-130.
2. R. Yang et al., "PixelFlex: A Reconfigurable Multi-Projector Display System," *Proc. IEEE Visualization*, IEEE CS Press, 2001, pp. 21-26.
3. R. Raskar et al., "Shader Lamps: Animating Real Objects with Image-Based Illumination," *Proc. Eurographics Rendering Workshop,* Eurographics Assoc., 2001, pp. 89-102.
4. A. Majumder et al., "Achieving Color Uniformity Across Multi-Projector Displays," *Proc. IEEE Visualization*, IEEE CS Press, 2000, pp. 117-124.
5. E. Bier et al., "Toolglasses and Magic Lenses: The See-Through Interface," *Proc. ACM Siggraph*, ACM Press, 1993, pp. 73-80.
6. T. Yoshida, C. Horii, and K. Sato, "A Virtual Color Reconstruction System for Real Heritage with Light Projection," *Proc. Virtual Systems and Multimedia*, IEEE CS Press, 2003, pp. 158-164.

Oliver Bimber is a junior professor for augmented reality at the Bauhaus University Weimar in Germany. His research interests include future display technologies, real-time rendering, optics, and advanced human–computer interfaces. He holds a PhD in computer science from the Technical University of Darmstadt.

Franz Coriand is a graduate student of media systems at the Bauhaus University Weimar. His research interests include augmented reality and virtual reality, as well as software engineering and Unix-based operating systems.

Alexander Kleppe is a graduate student of media systems at the Bauhaus University Weimar. His research interests include augmented reality and computer graphics, as well as input devices and interaction techniques for virtual environments.

Erich Bruns is a graduate student of media systems at the Bauhaus University Weimar. His research interests include computer graphics, augmented reality, and the psychology of perception, as well as the relation between anthropology and computer science.

Stefanie Zollmann is a graduate student of media systems at the Bauhaus University Weimar. Her research interests include virtual reality and augmented reality, as well as input and manipulation techniques for 3D environments.

Tobias Langlotz is a graduate student of media systems at the Bauhaus University Weimar. His research interests include augmented reality, virtual reality, and multiuser game engines.

Readers may contact Oliver Bimber at obimber@computer.org.

For further information on this or any other computing topic, please visit our Digital Library at http://www.computer.org/publications/dlib.