

Superimposing Dynamic Range

Oliver Bimber,* Bauhaus-University Weimar and Daisuke Iwai,† Osaka University

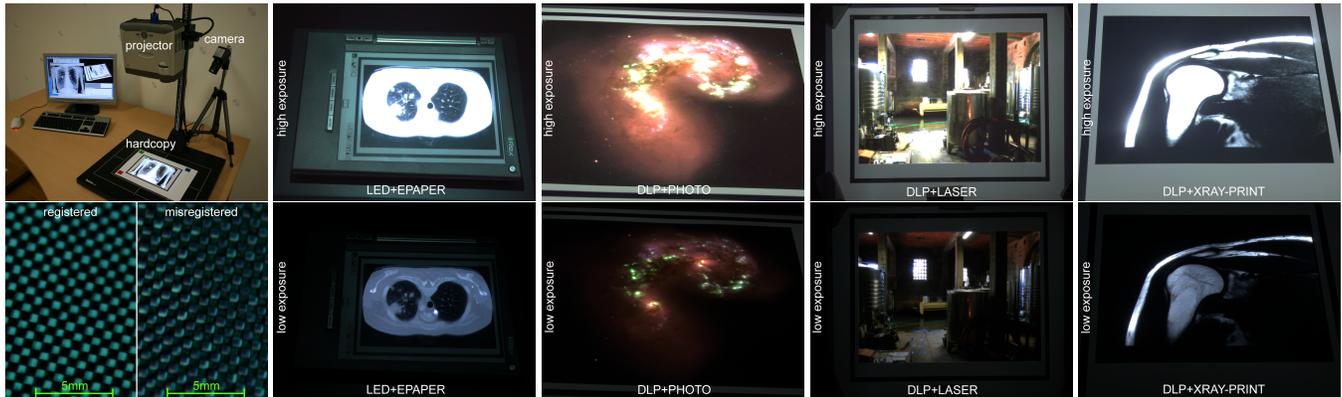


Figure 1: Registering a projector precisely to a paper print or to an interactive reflective display allows extending contrast, perceivable tonal resolution and color space beyond the capabilities of either printer, display or projector. From left to right: experimental setup and example for achieved registration precision (projected square wave grid on printed square wave grid with a field size of 0.6 mm), low and high exposure photographs of different hardcopies (ePaper display, photographic print, laser print and X-ray print) amplified with LED and DLP projectors.

Abstract

We present a simple and cost-efficient way of extending contrast, perceived tonal resolution, and color space of reflective media, such as paper prints, hardcopy photographs, or electronic paper displays. A calibrated projector-camera system is applied for automatic registration, radiometric scanning and superimposition. A second modulation of the projected light on the surface of such media results in a high dynamic range visualization. This holds application potential for a variety of domains, such as radiology, astronomy, optical microscopy, conservation and restoration of historic art, modern art and entertainment installations.

CR Categories: I.3.3 [COMPUTER GRAPHICS]: Picture/Image Generation—Display algorithms; I.4.0 [IMAGE PROCESSING AND COMPUTER VISION]: General—Image displays; I.4.1 [IMAGE PROCESSING AND COMPUTER VISION]: Digitization and Image Capture;

Keywords: Projector-camera system, hardcopy image, HDR display, luminance quantization, HDR splitting, inverse tone-mapping.

*bimber@uni-weimar.de

†daisuke.iwai@sys.es.osaka-u.ac.jp

1 Introduction and Motivation

Following pioneering work on recovering high dynamic range (HDR) radiance maps from photographs [Mann and Picard 1995; Debevec and Malik 1997], much research has been carried out since then on capturing HDR content as well as on displaying it on low dynamic range (LDR) displays. A large body of tone-mapping techniques, such as the photographic tone reproduction operator [Reinhard et al. 2002] and many others, aim at visually matching the appearance of a tone-mapped HDR image with the observed scene when being displayed on an LDR screen.

It was only recently that HDR displays were introduced which could present content over several orders of magnitude between minimum and maximum luminance. These displays are based on the principle of double modulation of light. For example, transmissive spatial light modulators, such as LCD panels, are coupled with a locally controllable back-illumination for achieving a high contrast.

Our Contribution. We present a simple and low-cost method of superimposing high dynamic range visualizations on arbitrary reflective media, such as photographs, radiological paper prints, electronic paper (ePaper), or even reflective three-dimensional items. Our technique is based on a secondary modulation of projected light being surface-reflected. This allows boosting contrast, perceivable tonal resolution, and color saturation beyond the possibility of projectors (when projecting onto regular screens), or the capability of spatially uniform environment light (when illuminating such media). Figure 1 illustrates various examples of this.

Besides the conceptual contribution, our technical contributions are the following: We explain how to boost the dynamic range of existing low dynamic range prints, such as ordinary photographic prints in photo albums or picture books (Section 3.1). This technique is not limited to two-dimensional surfaces, but can also be applied in combination with three-dimensional objects. We explain how existing HDR images can be split appropriately into projected images and printed images that together result in a correct HDR appearance (Section 3.2). Finally, we present a luminance quantization technique which maximizes the number of perceived tonal values while considering the discrete nature of the applied modulation devices,

such as projectors and printers (Section 3.3). This is essential for visualizing radiological data with as many differentiable gray levels as possible.

Preview of Results. In our experiments, we achieved physical contrast ratios of 45,000-60,000:1 with a peak luminance of more than $2,750 \text{ cd/m}^2$, and could technically reproduce more than 620 perceptually distinguishable tonal values when using the mapping function described in [Mantiuk et al. 2005; Mantiuk et al. 2004]. Furthermore, we attained color space extensions of up to a factor of 1.4 (compared to a regular projection on white screens) or factor of 3.3 (compared to regular paper prints under environment light). Thereby, the spatial image resolution of the prints can be several hundred lines per inch (*lpi*) while their raster resolution is normally several thousand dots per inch (*dpi*). We modulate luminance and chrominance with a registration error of less than 0.3 mm . Given this registration precision, our contrast frequency¹ can be up to 21 line pairs per inch (*lppi*) – or 7 cycles per degree (*cpd*) when viewed from a common distance of 50 cm . The automatic projector registration takes 5-15 seconds. Thereby, the reflective media can be flat or arbitrarily bent as long as its surface does not show discontinuities from the perspective of the projector or the camera. The rastered modulation outcome does not produce visible artifacts, such as moiré patterns, as would be the case for related transmissive modulation approaches.

Application Example. Specialized printing technologies, sometimes referred to as X-ray prints, are being used in the medical field for exchanging radiological images. Compared to conventional hard-copy media, such as X-ray film, they offer significant cost reductions, longer durability (they are less sensitive to light) and colored visualization usage. However, paper-prints do not provide the diagnostic quality of X-ray film when viewed under environment light. Under environment light conditions, paper-prints provide a contrast that is far less than 100:1 (see Table 1). The contrast of X-ray film with an optical density of $D=4$ is in the order of 10,000:1. A high contrast, contrast frequency and spatial image resolution, as well as the reproduction of a large number of perceivable tonal values and a high peak luminance are critically important for radiological visualizations, such as in mammography. These are requirements that cannot be met by most interactive HDR or LDR displays. The application of the techniques described in this paper, however, has the potential to achieve diagnostic quality with superimposed paper-prints at a fraction of the cost of X-ray film development. Our approach represents a cost-efficient add-on for such print technologies that allows exceeding the contrast of X-ray film by up to factor of 6.

2 Related Work

Two areas are closely related to our approach: HDR displays and projector-camera systems. They are discussed below.

HDR Display Approaches. Ledda et al. [2003] present a passive stereoscopic HDR viewer that applies two overlaid transparencies for each eye for luminance modulation and achieves a contrast ratio of 10,000:1. Based on the initial work of Seetzen et al. [2003], active HDR displays were described in [Seetzen et al. 2004] that modulate images displayed on an LCD panel with a locally varying background illumination. This is either produced by a lower-resolution LED panel, or by a higher-resolution DLP projector. A contrast ratio of over 50,000:1 together with a peak luminance of $2,700 \text{ cd/m}^2$ (for the projector-based backlight) and $8,500 \text{ cd/m}^2$ (for the LED-based backlight) were reported. Rosink et al. [2006] describe an HDR display prototype that utilizes two parallel-aligned (5 Mpixel) LCD panels which together enable a per-pixel contrast of

3,000–100,000:1 (depending on the viewing angle) together with a peak luminance of $1,000 \text{ cd/m}^2$. A color LCD backlight modulator behind a monochrome front panel allows for the adjustment of the exact white point of the display. Pavlovych et al. [2005] proposed an HDR projector that modulates the image path (i.e. after image generation) with a low-resolution monochrome LCD panel and a set of lenses placed in front of a regular DLP projector's object lens. It attains an ANSI-contrast of 708:1 and a peak luminance of 425 cd/m^2 . Kusakabe et al. [2006] describe an HDR projector that modulates the illumination path (i.e. between light source and image generator) with three (RGB) low-resolution LCoS panels, before a subsequent luminance modulation occurs with a high-resolution monochrome LCoS panel. Similar approaches are explained by Damberg et al. [2007] who achieve a contrast of 2,695:1. A variety of inverse tone-mapping techniques is now being developed, such as Banterle et al. [2006] and others, to convert existing LDR content into an HDR format to be viewed on such devices. All of these systems share three common properties which are essentially different from our approach:

Firstly, they apply mainly a second image modulation on a transmissive basis (either through transparencies or LCD/LCoS panels that utilize transmissive color and polarization filters) and consequently suffer from a relatively low light-throughput (e.g. regular color / monochrome LCD panels transmit less than 3-6% / 15-30% of light) and therefore require exceptionally bright light sources. A reflective secondary modulation of the image is more efficient. For example, approximately 70-90% of the irradiating light is reflected by standard paper.

Secondly, with Rosink et al. [2006] as an exception, one of the two modulation images is of low-resolution and blurred in order to avoid artifacts such as moiré patterns that are due to the misalignment of two raster modulators, as well as to realize acceptable frame-rates. Thus, high contrast values can only be achieved in a resolution of the low-frequency image. In our case, the two raster modulators (e.g. a printed image and a projected image) are always arbitrarily aligned. But since the raster resolution of printers is much higher than the one of projectors (up to a factor of 30), the Nyquist-Shannon theorem is always satisfied, and visible moiré patterns are not produced. Irregular dithering patterns are also beneficial for avoiding moiré artifacts. Consequently, a high contrast frequency is achieved, and neither the projected nor the printed image has to be blurred. In addition, a high spatial image resolution of currently up to 300 lpi is supported by state-of-the-art printing technology (up to 150 lpi for the devices used in our experiments). Thus we benefit from the modulation of two high-frequency images.

Thirdly, since one of the two images is monochrome (mainly to reach a high peak luminance), only luminance is modulated, while chrominance modulation for extending the color space is in some cases considered to be future work. We show that an extension of color space can in fact be achieved through the double modulation of chrominances. Thereby, the remaining white-light fraction of the projector's light bulb that passes its imperfect color filters is filtered again by being reflected on the print. This leads to more saturated colors.

Projector-Camera Approaches. A variety of projector-camera systems exist that perform radiometric or photometric compensations when projecting onto non-optimized (colored or textured) surfaces, such as in Grossberg et al. [2004] and others. In contrast to these techniques, our goal is to extend the dynamic range on reflective media rather than fitting arbitrary image content into the limited contrast range of non-optimized projection surfaces with varying reflectance properties. Majumder and Welch [2001] have shown that the intensity resolution (the number of gray scale levels per unit intensity) can be increased with multiple overlapping projectors. This leads to superior high-fidelity imagery. Yet, the contrast cannot be increased in this case.

¹The resolution at which a high contrast can be modulated.

3 Reflective Secondary Modulation

As described in Section 2, a secondary image modulation that is based on reflective filtering rather than on transmissive filtering has several advantages in terms of light-throughput and contrast frequency.

A more notable benefit of a reflective secondary modulation for high dynamic range visualization, however, is that the second modulation process can take place directly on arbitrary reflective surfaces that are not a necessarily a component of an HDR display. Examples include not only paper prints and photographs, but also historic and modern artwork and other physical two- or three-dimensional items. This section will explain how a reflective secondary modulation on physical surfaces can be achieved with calibrated projector-camera systems. The following techniques require that the geometric properties, as well as photometric response functions and transfer functions of camera, projector and printer (or surface) are known and can be calibrated. Furthermore, the reflective medium and the projector must be precisely registered. Note that although the following techniques are explained based on examples of static paper prints and photographs, they are equally applicable to interactive ePaper displays. The method described in Section 3.1 can also be used in combination with three-dimensional items.

3.1 Boosting Low Dynamic Range Photo Prints

In our first example we explain how to boost the dynamic range of existing low dynamic range prints. This could be, for instance, an ordinary photographic print in a photo album or in a picture book. We do not assume that corresponding high dynamic range representations exist, or that the prints are linearized.

3.1.1 Digital Reconstruction and Inverse Tone-Mapping

First, we measure the maximal contrast values that can be achieved locally as well as globally by reflecting projected light on the print. We achieved this by capturing two high-resolution HDR images of the print – one under a full white and under a full black projection. These two images represent the local maximum (I_{max}) and minimum (I_{min}) reflectance at every single point on the print in RGB space. Converting this data to local luminance values, the global maximum (\bar{L}_{max}) can be determined from I_{max} and the global minimum (\bar{L}_{min}) can be determined from I_{min} .

The next step is to derive an HDR representation (I_{target}) of the printed image content that we are targeting to visualize. For this, two possibilities exist: If the corresponding digital photograph of the print is available, we can apply an inverse tone-mapping operator to compute a possible HDR representation. If no corresponding digital photograph is available, but the transfer function of the printer is known, then we can reconstruct the original colors from the printed ones by applying the printer’s inverse transfer function to I_{max} . The result is then being used as input for an inverse tone-mapping operator. As mentioned above, this does require that the projector and the camera are photometrically calibrated and linearized.

Figure 2 illustrates two examples for reconstructing the original digital images from photographs of their photo prints. The chrominance deviation of the reconstructions is less than 4% (measured in CIE $L^*a^*b^*$ space) when comparing them to their originals.

In some of our experiments, we use the inverse tone-mapping operator described by Banterle et al. [2006], which reverses the global version of the photographic tone reproduction operator [Reinhard et al. 2002]. In [Akyüz et al. 2007], however, it was argued that from a perceptual point of view, simple scaling transformations can yield similar results as sophisticated tone-mapping operators (and sometimes even outperform them). Thus, we can also adapt their

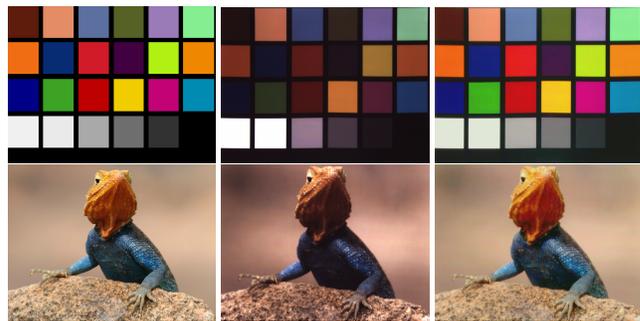


Figure 2: Color reconstruction of Macbeth samples (top) and digital image (bottom) developed on photographic paper with a Kodak System 88 digital lab system. The original (left), rectified camera image of photo without (center) and with (right) reconstructed colors.

scaling transformation for estimating an HDR representation:

$$L_{target} = (\bar{L}_{max} - \bar{L}_{min}) \left(\frac{L - \tilde{L}_{min}}{\tilde{L}_{max} - \tilde{L}_{min}} \right)^\gamma + \bar{L}_{min}, \quad (1)$$

whereby –as above– L are the luminance values of the reconstructed digital image, and \tilde{L}_{max} , \tilde{L}_{min} are the respective global maximum and minimum of L . After converting L_{target} to RGB (I_{target}) by recombining the new luminance values with the original chrominances, Equation 2 can be used for estimating the projection image. Once the HDR representation has been approximated, Equation 2 is used to compute the normalized projection image (I_P) based on the measured local contrast boundaries:

$$I_P = (I_{target} - I_{min}) / (I_{max} - I_{min}) \quad (2)$$

The modulation of I_P on the print will lead to the desired HDR appearance (I_{target}). The fourth column in Figure 1 shows an example for which this technique has been applied.

Note that since we do not assume a linear transfer function of the printer in this case, a minimum of clipping during compensation with Equation 2 is ensured if γ in Equation 1 equals the gamma of the hardcopy device’s transfer function.

3.1.2 Notes on Calibration and Registration

The above techniques require a precise geometric registration between projected, captured and printed images. However, we neither assume that the projector and camera are optically co-aligned, nor that the prints are perfectly flat. Furthermore, the photometric behaviour of all components has to be known. The notes below provide a guideline on how this can be achieved².

Projector-Camera-Print Registration. Assuming that the print is arbitrarily shaped, but does not contain geometric (but possibly radiometric) discontinuities, the pixel-correspondence between the projector and camera over the print surface can be determined efficiently through structured light techniques. However, care has to be taken that such a technique is robust, even for non-uniformly colored and dark surface portions that absorb a large amount of projected light. Furthermore, taking a minimum number of images speeds up the calibration process, but more importantly, it prevents placing too much stress on the mechanical parts in case digital SLR cameras are used for registration. Our scanning technique requires capturing only three images for achieving a robust projector-camera-surface registration.

²Details on the techniques that are outlined in this section, or in sections 3.2.2 and 3.3.2 can be found in the supplementary material.

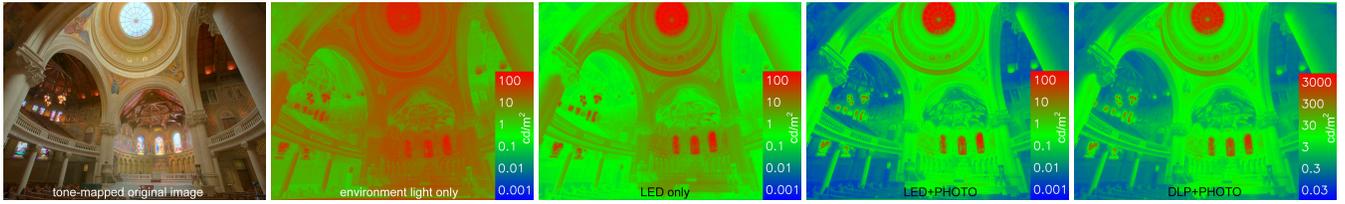


Figure 3: From left to right: Measured contrast of tone-mapped HDR image as photographic print under environment light and as projection on a white screen. Split HDR content modulated with LED+PHOTO and DLP+PHOTO – leading to contrast enhancements of 2-3 orders of magnitude.

Softcopy-Hardcopy Registration. If the registration of an existing digital image and the captured photograph of the corresponding print is required, we apply wide-baseline feature matching as explained in [Lepetit and Fua 2006]. All matched feature points are triangulated, and missing correspondences inside and outside the convex hull of the constructed triangle mesh are interpolated and extrapolated respectively. The resulting look-up table provides pixel correspondences between the printed image that are visible in the photograph and its original digital image. Note that such a registration is not necessary if I_{max} is used for estimating the HDR image, since it is already registered in the perspective of the camera.

Photometric Projector-Camera Calibration. Both the projector as well as the camera have to be linearized. Furthermore, the non-linear light drop-off of the projection, the contribution of the environment (including the projector’s black level) on the print, as well as the color mixing taking place between projector and camera, have to be measured and compensated for all projected and captured images. We apply standard calibration techniques for projector-camera systems, which are explained in detail in [Brown et al. 2005; Bimber et al. 2007]. If projectors with pulse-width modulation are applied (such as DLP or GLV), the projection’s refresh time must always be an integer multiple of the camera’s exposure time, so as to ensure the correct integration over all colors and intensities. For displaying purposes only, an arbitrary refresh rate can be chosen.

Measuring the Printer’s Transfer Function. Once the projector-camera system is fully calibrated as outlined above, the transfer function of the printer can be determined. We measure the full (color and intensity) transfer of the printer by printing and capturing all possible hues and tonal values. Thus, for an 8-bit RGB photo printer as an example (such as the Kodak System 88, which applies an 8-bit RGB LCD panel for light modulation during development), all 2^{24} values can be spatially encoded and printed on four letter-sized color charts. These charts are captured under a uniform white projector illumination. They are rectified and indexed, and their entries are sampled, smoothed, and stored in a look-up table. The inverse transfer function is simply the look-up table inverse. However, when reversing it, multiple entries must be averaged and missing values inside the convex hull of the sampled points must be interpolated. Missing values outside the convex hull are mapped to their closest valid entries. For higher precision, we store and apply the full look-up tables instead of separating individual color channels and fitting them into a set of analytical functions. Note that these look-up tables sample only the color and intensity transfer up to a scale.

3.2 Displaying High Dynamic Range Imagery

The next example explains how to visualize the correct appearance of an existing HDR image (I_{HDR}) through the modulation of the two LDR images (I_A, I_B). Thus, the HDR image has to be split appropriately into the corresponding projected image and printed image.

3.2.1 HDR Splitting

If both the projector and printer are linearized, such a splitting can generally be performed with

$$I_A = TM_{AB}(I_{HDR})^{\gamma \frac{a}{a+b}}, \quad (3)$$

$$I_B = TM_{AB}(I_{HDR})^{\gamma} / T_A(I_A), \quad (4)$$

where A is the device (i.e. either projector or printer) with a possibly significantly lower tonal or spatial image resolution compared to the other device (B). Artifacts, such as banding, that appear in image I_A due to a low quality of device A are compensated with image I_B . TM_{AB} is an initial tone-mapping operator that has to be applied if I_{HDR} exceeds the dynamic range that results from the modulation of A and B . We use linear tone-mapping if the final result must equal the original data up to a scale. An optional gamma correction can be applied before the mapped HDR content is split relative to the individual bit depths (a, b) that are supported by both devices. As for other tone-mapping operators, this allows the modification of the appearance of the final result based on user preferences. Again, γ has to be equal one if the final result must remain linear.

Thus, Equation 3 computes the image for the lower quality device with bit depth a , while Equation 4 compensates for artifacts in it using the higher quality devices with bit depth b . The bit depth of conventional projectors and printers is normally 8, while it is 4 for most commercial ePaper displays. T_A is the linearized transfer function of A that allows simulating the appearance of I_A with respect to A ’s tonal and spatial image resolution. Artifacts that result from the transfer function of the better device (B) cannot be compensated and will remain visible.

In principle, Equations 3 and 4 are generalizations of the splitting techniques described in [Seetzen et al. 2004; Trentacoste et al. 2007]. Instead of assuming two devices with equal bit depths, however, our technique supports devices with unequal bit depths. An example that was computed with this technique is shown in Figure 3.

3.2.2 Notes on Calibration and Registration

The geometric registration and the photometric calibration of projector, camera and printer for this case are the same as explained in Section 3.1.2.

Since the hardcopies are now custom-printed based on the individual splitting outcome, we can enhance the precision of softcopy-to-hardcopy registration by printing an additional registration frame. Although feature-based registration techniques are quite invariant to photometric differences between a captured hardcopy and digital softcopy, their precision depends very much on the number and on the distribution of detected feature points – and consequently on the image content. The additional registration frame enhances the result of this registration step.

3.3 Visualizing Radiological Data

In our third example, we explain how radiological images can be displayed effectively. In contrast to the techniques described above, it is now essential to produce a large number of perceptually distinguishable gray scales, rather than only a high physical contrast and dynamic range.

3.3.1 Luminance Quantization

As for other HDR display approaches, our reflective secondary modulation leads to a large number of physically producible luminance levels. Due to the non-linear response of the human visual system however, not all of them are perceptually distinguishable. The number of discernible luminance levels (*Just Noticeable Difference steps*, or *JND steps*) increases with a rising peak luminance of the display [Seetzen et al. 2004].

Since an exact representation with guaranteed distinguishable luminance levels is essential for several professional applications such as radiology, radiological images should be converted to a perceptually linear JND space, rather than being presented in a physically linear luminance space. All diagnostic monitors support this mapping through integrated look-up tables that are the result of frequent display calibrations.

Ghosh et al. [2005], for instance, describe perceptually linear transfer functions for volume rendering on HDR displays that modulate images on a transmissive basis. However, the technically achievable luminance space of displays that apply two image modulations (reflective or transmissive) is discretized and holds a quantization challenge. On the one hand, selected JND steps may not be achieved exactly, since they do not map to a producible luminance level. This is especially the case if both modulators are independently linearized (leading to a reduction of tonal values per se in each individual channel) or have a small local tonal resolution. Also many similar luminance levels can be approached with more than one modulation combination, as it is illustrated in Figure 4. These two issues have been ignored in the splitting technique explained in Section 3.2, or in related techniques, such as [Seetzen et al. 2004; Ghosh et al. 2005; Trentacoste et al. 2007]. They assume that each computed luminance level or JND step can be achieved by the two modulators, or that the nearest possible match is sufficient. This assumption is adequate for non-critical applications, but it will fall short of the high demand of radiological diagnostics. Not considering the discretization of the modulators can lead to indistinguishable luminance values and therefore to a smaller number of JND steps than can be produced.

This raises the following question: How can the selected JND steps be optimally mapped to individual modulator responses, in such a way that a maximum of JND steps is technically achieved, and that the combination of both modulators' transfer functions which produce the selected JND steps, remains as monotone as possible? The second condition is important to avoid visual artifacts in case of slight misregistrations, significant differences in modulator resolution, or imprecision occurring in their measured transfer functions. Intensity artifacts become clearly visible in these cases if the selected modulator responses that lead to similar luminance values differ significantly and are not precisely aligned, as shown in Figure 7 (bottom-right). For displaying gray scale content, we solved this problem by actually measuring the discrete luminance values that can technically be achieved through the modulation of different printed and projected gray scales, instead of ignoring the quantization problem when computing the modulator responses for linearized devices.

After having measured the possible luminance levels, we can color-code and plot them in a normalized gray scale space (e.g. with x for printer gray scales and with y for projector gray scales, with $0 \leq x, y \leq 1$), as shown in Figure 4.

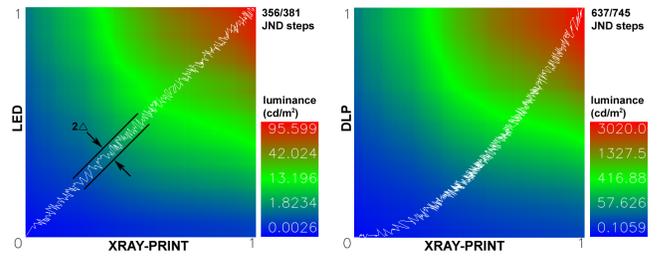


Figure 4: Luminance values in normalized gray scale space and selected JND steps (white) for combinations (see Section 4): LED+XRAY-PRINT and DLP+XRAY-PRINT with $\Delta = 0.025$.

Our goal is to fit a monotonic basis function to the measured luminance values that leads to an optimal assignment of different JND steps to unique gray scale pairs. To meet the two conditions described above, we fit the function $y = x^\sigma$ with the following objective:

$$\max\left(\left|\bigcup_{j=0}^n \{s_j \mid \min_{s_j \in C_j} (L_{s_j} - L_j)\}\right|\right), \quad (5)$$

$$C_j = \{c \mid \Delta_c < \Delta, L_j \leq L_c\},$$

where $j = 0..n$ indexes each individual JND step (with corresponding luminance L_j) that can be derived from lowest reflected valley luminance and highest reflected peak luminance - which are both known after calibration.

With these, we apply the mapping function described in [Mantiuk et al. 2005; Mantiuk et al. 2004] for assigning JND steps to luminance values, since it is defined for our luminance range. Thereby, L_0 is equivalent to the lowest black level reflection of the projector. For each theoretically possible JND step (j) with luminance L_j we choose a set (C_j) of gray scale candidates ($c \in C_j$) that leads to reproducible luminance levels (L_c) larger than or equal to L_j , and whose shortest (x, y)-distance (Δ_c) to our basis function is not larger than a predefined maximum (Δ). From each C_j , we select the candidate $s_j \in C_j$ that is closest to L_j .

Fitting our basis function while maximizing the number of technically achievable JND steps, results in one optimal set of projector and printer gray scales for each JND step that satisfies our conditions. These are the gray scales that belong to the selected samples s_j for each JND step j at the optimal σ . Note that we chose the basis function $y = x^\sigma$ for curve fitting, because it can address the entire gray scale space while optimizing only a single parameter.

The gray scale selections that are plotted in Figure 4 illustrate results found with our quantization technique for different projectors. They show that the optimal solution is not necessarily at $\sigma=1$. The two images at the bottom-right of Figure 7 show the reduction of registration artifacts for a modulation on an ePaper display when curve-fitting is more constrained through a smaller Δ value. This, however, will also lead to a slightly lower number of achievable JNDs steps. Note that Δ is currently empirically set by the user. Deriving it automatically from perception constraints belongs to our future work.

For displaying color content, the luminance of the original RGB values are scaled with the corresponding (normalized) gray scales that have been selected for printer and for projector. Thus, the luminance variation of the resulting two color images that are modulated is the same as is the case when displaying gray scale content (cf. Figure 7–bottom row for an example that modulates an 8-bit color projection on a 4-bit gray scale ePaper display).

3.3.2 Notes on Calibration and Registration

In contrast to the calibration requirements that are explained in Sections 3.1.2 and 3.2.2, neither projector nor printer are linearized in this case. For measuring absolute luminance and chrominance values rather than device dependent RGB values, we initially match the camera response with corresponding values delivered by a spectroradiometer.

Measuring Discrete Luminance Space. If the printer transfer function is measured as explained in Section 3.1.2, and converted from RGB values to absolute luminance and chrominance values, then the reflected luminance values (up to scale) of all 2^h printed gray scales can be directly selected from it. The reflected luminance of all 2^p projector gray scales are measured by projecting them onto a flat, white hardcopy sample. Their multiplication leads to the corresponding luminance values for all 2^{h+p} gray scale combinations.

Projector-Camera Registration. For professional applications, such as the visualization of radiological images, we can ensure that the prints remain planar. Technically, we achieve this by using a custom-built table that straightens the prints through an electrostatic force. In case of planar surfaces, projector-camera registration then boils down to a simple homography estimation. The 3x3 homography matrix that correlates the projector and camera pixels over the table-top surface has to be initially calibrated, which is then independent to the actual position of the print on the table.

Softcopy-Hardcopy Registration. The registration of prints located on the table plane with correlated camera pixels is simple when using a registration frame around the printed image content. Since the projector-camera homography matrix remains constant, new prints can be re-registered quickly by detecting their registration frame through a single photograph.

4 Results

On average, the registration error between the projected image and printed image (or ePaper display) is less than 0.3 mm (cf. Figure 1 bottom-left) for the devices applied in our experiments. No visible moiré patterns are produced through double modulation, which can be contributed to the high raster resolution and irregular dithering patterns of the printers and the ePaper display. The remainder of this section presents other quantitative measurement results for different modulator combinations, as well as an initial informal and subjective feedback of professional users.

4.1 Quantitative Measurements

Table 1 presents measurement results for the following projectors, printers and ePaper display:

A Kodak System 88 professional digital lab system for photographs (*PHOTO*) together with Kodak Royal N photo paper, a 2,400 dpi Xerox WorkCentre 7655 professional printer used for medical data (*XRAY-PRINT*), a 1,200 dpi Samsung CLP-510N consumer color laser printer (*LASER*), an XVG A iRex iLiad ePaper display (*EPAPER*), an SVGA MITSUBISHI PK20 LED projector (*LED*), and a XGA Optoma DX733 DLP projector (*DLP*). For taking measurements under environment light (*ENVIRONMENT*), a 20W halogen lamp was applied, while for regular screen projections (*SCREEN*) a blank sheet of regular laser printer paper was used.

We alternatively apply one of the two CMOS cameras for calibration and measurements: a 10.1Mpixel Canon EOS Digital Rebel XTi consumer digital SLR camera, and a QXGA ARTRAY ARTCAM 300MI professional image processing camera. Both have been linearized and calibrated to absolute values with a spectroradiometer. While the ARTCAM is faster, the EOS offers a higher resolution and

	parameters	PHOTO	XRAY-PRINT	LASER	EPAPER	SCREEN
DLP	CIE xy coverage (%)	36.4	38.1	37.5	n.a.	26.8
	contrast ratio	46,679:1	28,516:1	15,595:1	6,458:1	654:1
	peak luminance (cd/m ²)	2,750	3,020	3,100	1,440	3,100
	black level (cd/m ²)	0.0601	0.1059	0.1988	0.2230	4.7401
	#JND / #JND max	628/743	637/745	623/738	329/644	209/667
LED	CIE xy coverage (%)	58.6	59.4	59.2	n.a.	55.7
	contrast ratio	61,854:1	36,167:1	23,618:1	8,458:1	945:1
	peak luminance (cd/m ²)	87.1	95.6	97.3	43.2	97.3
	black level (cd/m ²)	0.0014	0.0026	0.0041	0.0051	0.103
	#JND / #JND max	345/373	356/381	349/380	201/286	169/335
ENVIRONMENT	CIE xy coverage (%)	10.9	10.5	9.8	n.a.	n.a.
	contrast ratio	73.3:1	40.4:1	32.2:1	9.79:1	n.a.
	peak luminance (cd/m ²)	67.1	70.3	77.2	37.7	77.2
	black level (cd/m ²)	0.9154	1.7401	2.3975	3.8509	n.a.
	#JND / #JND max	128/255	126/247	118/250	16/155	n.a.

Table 1: Results on different combinations of projectors, printers and ePaper display.

more f -stops. Our experimental hardware configuration is shown in Figure 1 (top-left).

In most combinations, contrast ratio, color space coverage, and number of technically achieved JND steps (maximized #JND by considering modulation discretization, as explained in Section 3.3.1) versus theoretically possible JND steps (computed #JNDmax – from valley and peak luminance only, as described in [Mantiuk et al. 2005; Mantiuk et al. 2004]) are boosted significantly. The extension of the color space (cf. Figure 5) results from the fact that the remaining white-light portion which is transmitted by the imperfect color filters of the projectors is additionally filtered when being reflected by the hardcopy. In general, this leads to more saturated colors.

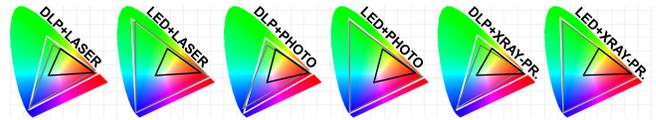


Figure 5: Coverage of CIE xy chromaticity space measured with a spectroradiometer: prints under environment light (black), projectors on white screen (gray), and combinations (white).

The LEDs of the LED projector, however, are already quite saturated, which leads to only minor extensions of the color space. They are also relatively dim, and a large number of JND steps are not possible due to the low peak luminance of the LEDs. However, the largest overall contrast ratio can be achieved. Due to the low native contrast and tonal resolution (only 16 gray scales) of the ePaper display, the overall contrast and number of JND steps that were achieved with it were roughly less than half compared to the other combinations. Yet, ePaper allows for displaying interactive content.

4.2 Subjective Professional Feedback

We presented our *XRAY-PRINT+DLP* prototype together with radiological datasets (one thorax CR scan, and one thorax CT scan with four different density settings and cutting planes – all monochrome) to ten professional radiologists. The images were visualized with our luminance quantization technique, as explained in Section 3.3.1. The radiologists were employed by different institutions, and were questioned independently. With their experience, we asked them to compare the image quality of our approach to the image quality of X-ray film and high contrast medical monitors. A direct side-by-side comparison of the same image content displayed with the different media, however, was not yet performed. As shown in Figure 6, the subjective impression of the professionals indicates that our approach performs significantly better in all categories. The chart presents the average scores and the ranges of variation. A formal clinical study will be carried out in future. This early informal feedback, together with our quantitative measurements, however,

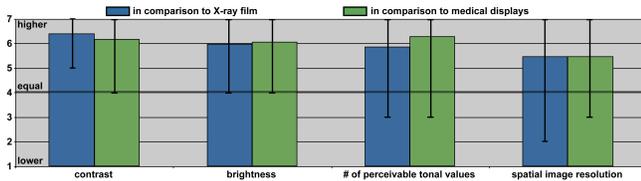


Figure 6: Questionnaire results from ten professional radiologists.

underlines our initial statement that a reflective modulation on radiological paper prints has indeed the potential to achieve the image quality of professional X-ray film.

5 Discussion

Our main goal with the presented approach is to enable high contrast visualizations on reflective media, such as paper prints, photographs, or even three-dimensional surfaces. It holds potential in several domains, such as radiology, optical microscopy, conservation and restoration of historic art, astronomy, modern art, entertainment, and more. Our goal however is not to compete with interactive displays. Rather, our proposed technique is complementary to applications that operate with high quality hardcopy images or other physical items. Yet, ePaper also allows for interactive presentations. This section summarizes the benefits and the limitations of our approach.

5.1 Benefits

High Contrast Frequency. Compared with most existing interactive HDR displays, our approach allows displaying images with a high spatial image resolution as well as with a high contrast frequency. Given our current registration precision of 0.3 mm , we can ensure a different modulation of two adjacent points at a distance of 0.6 mm (cf. Figure 1 bottom-left). Thus, we achieve a contrast frequency of 21 lppi ($=42\text{ lpi}$, or 7 cpd when viewed from a common distance of 50 cm). In addition, the spatial image resolution can be 150 lpi (e.g. for a raster resolution of $2,400\text{ dpi}$ and 256 printed gray scales³) and more. Thus our approach is well suited for near distance viewing of high quality static image content.

Efficient Light Throughput. A high contrast frequency and a high spatial image resolution are important for radiological visualizations. Therefore, initial displays that achieve a pixel-accurate backlight modulation are being developed [Rosink et al. 2006]. However, a double modulation through transmissive layers, such as LCD panels, will always suffer from extremely low light throughput, and an adequately high peak luminance that produces many distinguishable tonal values may be difficult to achieve. A reflective modulation, such as one in our case, is more efficient. It is also important to ensure that computed JND steps can truly be achieved, which is not necessarily the case due to a discretization of the modulators' responses. The high light throughput, and consequently the high brightness and the large number of guaranteed JND steps is another advantage of our approach. In addition, if LCDs are applied as modulators, the contrast is reduced depending on the viewing angle.

Extended Color Space. The double modulation of chrominances leads to a reduction of the remaining white-light fraction that pass the imperfect color filters of the projector. Since this extends the displayable color space, reflected colors appear more saturated.

³Dor digital halftoning, the spatial resolution of a printed image (in lines per inch – lpi) is normally the raster resolution (in dots per inch – dpi) divided by the square-root of the printed gray levels.

HDR on Arbitrary Reflective Surfaces. The second image modulation process can take place directly on arbitrary reflective surfaces that are not a component of an HDR display. This not only include two-dimensional surfaces, such as paper prints or ePaper, but is also easily amenable to three-dimensional surfaces. For example, operation microscopes that are equipped with projector-camera systems which are co-aligned with microscope's optics can boost the perceived contrast for surgeons in real-time. An optical contrast enhancement can also support conservators when restoring paintings, frescos, historic prints or artwork. Three-dimensional artistic or entertainment installation can be presented in brilliant colors and with a high contrast.

5.2 Limitations

Reduction of Physical Contrast. Environment light physically reduces the contrast. For this reason, high-contrast medical displays are only operated in darkened rooms. For shiny surfaces, such as glossy paper, specular reflections can reduce the contrast depending on the viewing angle in our case. For matte surfaces, such as regular paper prints or matte photo paper, the reflected light is almost evenly diffused.

Reduction of Perceived Contrast. The scattering of light in the eye reduced the perceived local contrast. This depends mainly on the adaptation luminance and on the spatial frequency of the observed content. Referring to the optical transfer function of the eye described in [Deeley et al. 1991] in combination with the relation of pupil diameter to known adaptation luminance in [Moon and Spencer 1944], we can still perceive a local contrast of 40%-69% of the global contrast at a contrast frequency of 21 lppi and an adaptation luminance of $0.06\text{ cd/m}^2 - 2,750\text{ cd/m}^2$. It is also reported in [Mantiuk et al. 2005] that the contrast sensitivity (CS) for an adaptation luminance above $1,000\text{ cd/m}^2$ is maximal at this contrast frequency (i.e. 7 cpd), and that the CS-peak only shifts to lower frequencies with a decreasing adaptation luminance.

Registration and Calibration Precision. The precision of geometric registration and photometric calibration strongly depends on the resolution and dynamic range of camera and projector. Strong highlights reflected towards the camera will degrade the overall calibration outcome. For non-planar surfaces, the depth-of-fields of the projector and camera are also an essential factor for avoiding regional defocus blur (captured and projected). Choosing small aperture settings increases the depth-of-field of the camera at the cost of longer exposure times.

Acknowledgements

We thank Francesco Banterle of University of Warwick, Rafal Mantiuk of Max-Planck-Institute for Computer Science, Shinsaku Hiura of Osaka University, the National Cancer Institute, the European Space Agency and Aycan Ltd., for providing content and support. Thanks to all SIGGRAPH reviewers for valuable comments and suggestions. Credits – shoulder: <http://barre.nom.fr/medical/samples/>, nebular: <http://www.spacetelescope.org/images/index.html>, and chest: <http://www.cancer.gov/>.

References

- AKYÜZ, A. O., FLEMING, R., RIECKE, B. E., REINHARD, E., AND BÜLTHOFF, H. H. 2007. Do hdr displays support ldr content?: a psychophysical evaluation. In *Proc. ACM Siggraph*, vol. 26, 38.2–38.7.
- BANTERLE, F., LEDDA, P., DEBATTISTA, K., AND CHALMERS, A. 2006. Inverse tone mapping. In *Proc. Conference on Computer*

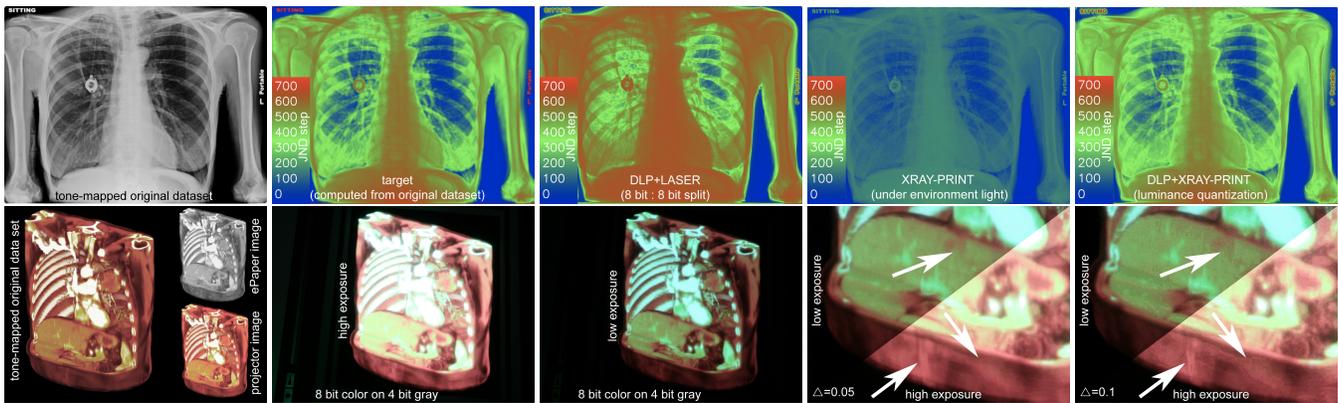


Figure 7: Top row: Comparing simulated target JND steps with measured JND steps of dissimilar modulation combinations (photographs rectified). Splitting leads to perceptually non-linear results, while plain X-ray prints support only a small number of JND steps. Bottom row: LED+EPAPER color on gray – revealing artifacts in close-up photographs (bottom-right) if a Δ that is too high is chosen for luminance quantization.

- Graphics and Interactive Techniques in Australasia and Southeast Asia, 349–356.
- BIMBER, O., IWAI, D., WETZSTEIN, G., AND GRUNDHOEFER, A. 2007. The Visual Computing of Projector-Camera Systems. In *Proc. Eurographics (State-of-the-Art Report)*, 23–46.
- BROWN, M., MAJUMDER, A., AND YANG, R. 2005. Camera Based Calibration Techniques for Seamless Multi-Projector Displays. *IEEE Trans. on Visualization and Computer Graphics* 11, 2, 193–206.
- DAMBERG, G., SEETZEN, H., WARD, G., HEIDRICH, W., AND WHITEHEAD, L. 2007. High dynamic range projection systems. In *Society of Information Displays Symposium Digest*, vol. 38, 4–7.
- DEBEVEC, P. E., AND MALIK, J. 1997. Recovering high dynamic range radiance maps from photographs. In *Proc. ACM Siggraph*, 369–378.
- DEELEY, R., DRASDO, N., AND CHARMAN, W. N. 1991. A simple parametric model of the human ocular modulation transfer function. *Ophthalmology and Physiological Optics* 11, 9193.
- GHOSH, A., TRENTACOSTE, M., AND HEIDRICH, W. 2005. Volume rendering for high dynamic range displays. In *Proc. EG/IEEE VGTC Workshop on Volume Graphics*, 91–231.
- GROSSBERG, M., PERI, H., NAYAR, S., AND BELHUMEUR, P. 2004. Making One Object Look Like Another: Controlling Appearance using a Projector-Camera System. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, vol. I, 452–459.
- KUSAKABE, Y., KANAZAWA, M., NOJIRI, Y., FURUYA, M., AND YOSHIMURA, M. 2006. YC-separation Type Projector with Double Modulation. In *Proc. International Display Workshop*, 1959–1962.
- LEDDA, P., WARD, G., AND CHALMERS, A. 2003. A wide field, high dynamic range, stereographic viewer. In *Proc. Conference on Computer Graphics and Interactive Techniques in Australasia and Southeast Asia*, 237–244.
- LEPETIT, V., AND FUA, P. 2006. Keypoint recognition using randomized trees. *IEEE Trans. on Pattern Analysis Machine Intelligence* 28, 9, 1465–1479.
- MAJUMDER, A., AND WELCH, G. 2001. Computer graphics optique: Optical superposition of projected computer graphics. In *Proc. Fifth Immersive Projection Technology Workshop*.
- MANN, S., AND PICARD, R. 1995. Being 'undigital' with digital cameras: Extending dynamic range by combining differently exposed pictures. In *Proc. IS&T 46th annual conference*, 422–428.
- MANTIUK, R., KRAWCZYK, G., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2004. Perception-motivated high dynamic range video encoding. In *Proc. ACM Siggraph*, vol. 23, 733–741.
- MANTIUK, R., DALY, S. J., MYSZKOWSKI, K., AND SEIDEL, H.-P. 2005. Predicting visible differences in high dynamic range images - model and its calibration. In *Proc. IS&T/SPIE's Annual Symposium on Electronic Imaging*, vol. 5666, 204–214.
- MOON, P., AND SPENCER, D. 1944. On the stiles-crawford effect. *Journal of the Optical Society of America* 34, 319–329.
- PAVLOVYCH, A., AND STUERZLINGER, W. 2005. A High-Dynamic Range Projection System. In *Proc. SPIE annual meeting*, vol. 5969.
- REINHARD, E., STARK, M., SHIRLEY, P., AND FERWERDA, J. 2002. Photographic tone reproduction for digital images. In *Proc. ACM Siggraph*, vol. 21, 267–276.
- ROSINK, J., CHESTAKOV, D., RAJAE-JOORDENS, R., ALBANI, L., ARENDS, M., AND HEETEN, G. 2006. Innovative lcd displays solutions for diagnostic image accuracy. In *Proc. Radiological Society of North America annual meeting*, electronic abstract.
- SEETZEN, H., WHITEHEAD, L. A., AND WARD, G. 2003. A high dynamic range display using low and high resolution modulators. In *Proc. Society for Information Display*, 1450–1453.
- SEETZEN, H., HEIDRICH, W., STUERZLINGER, W., WARD, G., WHITEHEAD, L., TRENTACOSTE, M., GHOSH, A., AND VOROZCOVS, A. 2004. High dynamic range display systems. In *Proc. ACM Siggraph*, 760–768.
- TRENTACOSTE, M., HEIDRICH, W., WHITEHEAD, L., SEETZEN, H., AND WARD, G. 2007. Photometric image processing for high dynamic range displays. *J. Visual Communication and Image Representation* 18, 5, 439–451.