Simulating and compensating changes in appearance between day and night vision

Robert Wanat Rafał K. Mantiuk*

Bangor University

Figure 1: *Retargeting from and to a dark display. Left: Image as seen on a 2 cd/m^2 peak luminance display. Center: Original image. Right: Bright image compensated for a 2 cd/m^2 display. When the original image is seen through a neutral density filter reducing luminance 100 times (2.0 D), it will match the appearance of the left image. When the right image is seen through the same filter thus simulating a dark display, it will appear similar to the original. Note that the seemingly exaggerated sharpness, color shift and brightness change are not perceived as such at low luminance levels. The images are best seen when the page is enlarged to $3/4$th of the screen width and viewed from about 0.5 m for a 24" monitor.*

Abstract

The same physical scene seen in bright sunlight and in dusky conditions does not appear identical to the human eye. Similarly, images shown on an 8000 cd/m^2 high-dynamic-range (HDR) display and in a 50 cd/m^2 peak luminance cinema screen also differ significantly in their appearance. We propose a luminance retargeting method that alters the perceived contrast and colors of an image to match the appearance under different luminance levels. The method relies on psychophysical models of matching contrast, models of rod-contribution to vision, and our own measurements. The retargeting involves finding an optimal tone-curve, spatial contrast processing, and modeling of hue and saturation shifts. This lets us reliably simulate night vision in bright conditions, or compensate for a bright image shown on a darker display so that it reveals details and colors that would otherwise be invisible.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Display Algorithms; I.4.2 [Image Processing and Computer Vision]: Enhancement—Greyscale manipulation, sharpening and deblurring

Keywords: contrast perception, color perception, retargeting, low luminance display, backlight dimming, computational display

Links: DL PDF WEB

*e-mails: {eepe11, mantiuk}@bangor.ac.uk

1 Introduction

Color and contrast perception varies significantly across the range of illumination levels. The most dramatic change in vision is observed when luminance drops below 3–5 cd/m^2, when the retinal cone cells steadily lose their sensitivity and visual signal is influenced by the retinal rod cells. In this, so called, mesopic vision range, we can observe a gradual loss of acuity and color vision. This important characteristic of the visual system is rarely taken into account when reproducing colors on electronic displays. While the state-of-the-art display colorimetry is almost entirely based on the cone-mediated vision (CIE 1931 color matching functions), a significant portion of the color gamut in modern displays often lies in the luminance range below 3 cd/m^2, partly mediated by rods. This is especially relevant for mobile displays, which can decrease their brightness down to 10-30 cd/m^2 of the peak luminance to reduce power consumption. This means that in the case of a 1000:1 contrast display that is dimmed, about 3/4th of the perceived color gamut cannot be accurately reproduced using traditional cone-based colorimetry.

To simulate and compensate the changes in image appearance due to a lower absolute luminance level, we propose a new appearance matching model and luminance retargeting method. The main application of the method is to compensate for appearance change in dimmed displays. When watching a display under low ambient illumination, it might be preferable to dim the backlight to reduce power consumption or avoid eye fatigue. In other situations, it could be desirable to reduce the backlight to preserve observer's dark adaptation, for example in a car mounted displays at night. The method can also perform retargeting in the opposite direction, from dark scenes to much brighter displays, in order to reproduce the appearance of night scenes, for example in driving simulators or in computer games.

The method relies on a model of color and contrast perception across the entire range of luminance. The change of overall brightness and contrast is compensated by optimizing the shape of a tone-curve so that it provides the best compromise between retaining

contrast and brightness while utilizing the available dynamic range of a display. The change in detail visibility is modelled using a novel suprathreshold contrast matching model. The change of color appearance is accounted for by modelling rod contributions and loss of color saturation at low luminance levels. All these components provide superior appearance matching across a range of luminance, which cannot be achieved with the existing color appearance models and tone-mapping operators.

The paper is organized as follows: after discussing related work, we describe in Section 3 the experimental methods used throughout the derivation. The novel method is formulated in Section 4 and its applications are presented in Section 5. A comparison with existing methods is made in Section 6 and the limitations are discussed in Section 7.

2 Related work

The **color appearance** community has recognized a number of appearance phenomena that are affected by absolute luminance levels [Fairchild 2005, ch.6]. The Hunt effect causes perceived colorfulness to increase with increasing surround and adapting luminance. The Helmholtz-Kohlrausch effect causes more saturated colors to appear brighter. The Stevens effect causes the range of perceived brightness levels to increase with higher adapting luminance. Bartleson and Breneman [1967] measured and modelled a similar effect of luminance and contrast, but for complex images and for varying levels of background luminance. They found the contrast of images to decrease with lower luminance. This is consistent with a common practice of reproducing images with a higher contrast (power exponent approx. 1.5) when projecting images in dark rooms, or the reproduction of television images (power exponent 1.25), which are typically watched in dim surroundings [Fairchild 2005, p.125].

Both Stevens and Hunt effects are included in the CIECAM02 color appearance model [Moroney et al. 2002]. If such an appearance model could accurately predict perceptual correlates, it could potentially be used to retarget content for different luminance levels. However, the existing appearance models suffer from several limitations: they predict color appearance only for photopic conditions, that is luminance levels above about 3 cd/m^2; they predict appearance for uniform color patches rather than complex images; and they do not consider the influence of spatial contrast on appearance, which, as we show later, is a crucial factor that is influenced by the absolute luminance levels. There have been attempts to extend appearance models to complex images [Kuang et al. 2007; Reinhard et al. 2012], but these methods were not intended to predict appearance at low luminance levels and were mostly intended for tone-mapping of high dynamic range images. In particular, the iCAM06 appearance model [Kuang et al. 2007] does not reproduce the appearance of night scenes [Petit and Mantiuk 2013]. Nevertheless, we evaluate the method of Reinhard et al. [2012] in Section 6.

Tone-mapping of high-dynamic-range images often involves reproducing the appearance of a scene under particular viewing conditions on a display offering a limited contrast and luminance range. Pattanaik et al. [1998] proposed one of the most comprehensive visual models, combining spatial vision, light and dark adaptation and appearance modeling. The model can simulate the appearance of a night scene on a regular display. Thompson et al. [2002] proposed a method simulating neural noise and the loss of acuity, which could be observed in low light conditions. Kirk et al. [2011] proposed a tone-mapping focused on color perception at low light conditions. A tone-mapping operator that optimizes a tone-curve for the target display and contrast reproduction was proposed in [Mantiuk et al. 2008]. We evaluate these methods in Section 6 and demonstrate that they neither provide sufficient accuracy nor account for all relevant effects.

Display backlight dimming. Many of consumer displays offer a feature to reduce the backlight when the ambient light is dim, mostly to lower energy consumption, but also to avoid eye fatigue due to high brightness. Some of these methods involve compensating for the dimmer backlight with increased transparency of the liquid crystal (LC) panel, so that the differences between the original image and the image shown with dimmed backlight is minimal [Chang et al. 2004]. The compensation methods may account for both temporal aspects to reduce flicker visibility [Iranli et al. 2006], and spatial aspects to reduce contours due to hard clipping [Kerofsky and Daly 2006]. None of these methods, however, compensates for the contrast perception change at low light. Our algorithm is directly applicable to the backlight dimming scenario.

Display color compensation. A method for adapting display color for low luminance conditions, which shares the same goals as our work, can be found in the patent literature [Kane and Kurtz 2012]. The method transforms an image using forward and inverse color appearance models (CIECAM97c) to retarget color appearance to lower luminance. This approach, however, assumes that the appearance model can correctly predict the changes in appearance between photopic and mesopic conditions for complex images. This is not the case for most appearance models, including CIECAM97c, as discussed in the color appearance paragraph above.

3 Experimental methods

The derivation of our algorithm was driven, calibrated and tested with strict experimental procedures to ensure a good appearance match between luminance levels. It is important to note that we did not assume correctness of the visual models from the literature, which were measured for simple stimuli. Instead, we tested them with complex images across a range of conditions. We found that a haploscopic matching method, where each eye is adapted to a different luminance level, gave the most consistent and repeatable results and therefore was used in all our experiments.

Images were shown on a colorimetrically calibrated 24" 1920×1200 LCD display (NEC PA241W) and viewed in a dark room. The display was driven with 10 bits per color channel and used the native extended color gamut. A piece of black cardboard was used to separate the display screen into two halves, so that each eye could see only one half of the screen. The viewing distance was restricted to 85 cm and the pixel angular resolution was 56 pixels per degree. Observers wore modified welding goggles, in which we removed the protective filter for one eye and introduced a photographic neutral density (ND) filter (Kodak Wratten 96 1D and 2D) for the other eye. The choice of the eye to cover with the filter was randomized between the sessions. Such a setup ensured that both eyes were adapted separately to different luminance levels and the visual glare did not affect the "darker" eye.

The observers were asked to adjust the parameters or make judgements so that the displayed image shown to the "dark" eye was as close as possible to the reference image shown to the "bright" eye (the method-of-adjustment). Each parameter adjustment session was completed by at least three expert observers for 10 images from the Kodak database[1] and the results were averaged. The observers were excluded from the comparison experiment described in Section 6.1.

[1] Kodak True Color Image Suite, available at `http://r0k.us/graphics/kodak/`

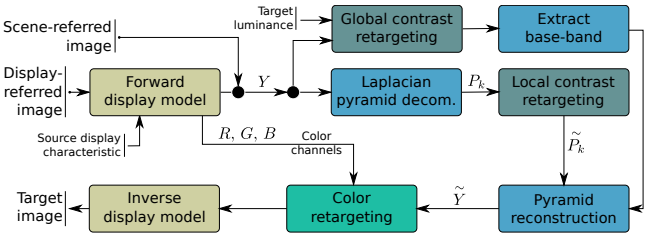

Figure 2: *The diagram of the proposed retargeting algorithm.*

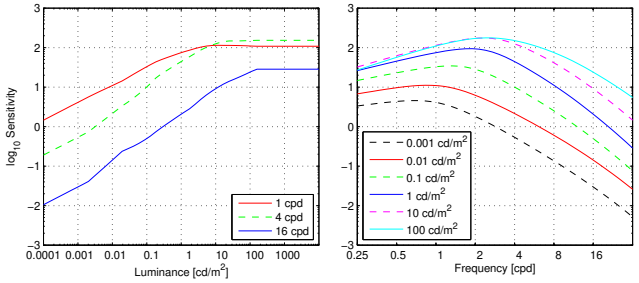

Figure 3: *Contrast Sensitivity Function (CSF) and its variation with luminance (left) and spatial frequency (right). The function is based on the model from [Mantiuk et al. 2011]. The frequency is given in cycles per degree (cpd).*

4 Algorithm and visual models

To model complex interaction between the absolute luminance level and image appearance, we analyze the problem with respect to three different aspects of an image: global contrast (tone curve), local contrast (detail) and color.

As shown in Figure 2, input to our algorithm is either a scene-referred image (a high dynamic range image, represented in absolute units), or a display referred-image, for example in the sRGB color space. In the latter case, the image needs to be transformed from gamma-corrected pixel values into absolute linear RGB values using a display model, such as a gamma-offset-gain (GOG) model [Berns 1996]. The GOG model should account for the effect of ambient light on the black level, similarly as done by Mantiuk et al. [2008]. Next, in the *global contrast retargeting* step the luminance channel of the image is modified by a global tone curve, which minimizes perceived contrast distortions when the image is shown on the target display. Only lower frequencies are extracted from that luminance image (*extract base-band* step) and the remaining spatial frequencies are retargeted separately for each frequency band of the Laplacian pyramid. The color changes caused by the influence of the rod vision are estimated based on the input and output absolute luminance levels and the new linear RGB values are calculated in the *color retargeting* step. Finally, the retargeted result of our algorithm needs to be transformed to pixel values using an inverse display model, or alternatively, into the sRGB color space. The following sections discuss each of these steps in detail.

4.1 Contrast retargeting

Before discussing contrast matching models, let us introduce two measures of contrast that we will use in this section. The Michelson contrast is defined as:

$$M = \frac{L_{max} - L_{min}}{L_{max} + L_{min}} = \frac{\Delta L}{L_{mean}}, \tag{1}$$

where L_{max} and L_{min} are the maximum and minimum luminance values of a sine wave, or alternatively ΔL is the modulation and L_{mean} is the mean value of a sine wave. The Michelson contrast varies between 0 and 1. When calculating image contrast in a multiscale representation, such as the Laplacian pyramid, it is more convenient to use the logarithmic contrast:

$$G = \frac{1}{2} \log_{10} \left(\frac{L_{max}}{L_{min}} \right). \tag{2}$$

The logarithmic contrast can be interpreted as the modulation of the sine wave in the logarithmic domain. We will use G and M symbols in the rest of the paper to distinguish between both measures. The following equations convert from one contrast to another:

$$G(M) = \frac{1}{2} \log_{10} \left(\frac{M+1}{1-M} \right), \quad M(G) = \frac{10^{2\,G} - 1}{10^{2\,G} + 1}. \tag{3}$$

Our ability to see small contrast (sensitivity) varies greatly with both frequency of the stimulus and its luminance. Such variations are well described by a number of Contrast Sensitivity Functions (CSFs) [Barten 1999], such as the one shown in Figure 3. The plots show the variation in sensitivity, which is the inverse of the threshold detection contrast. Although the CSF captures an essential characteristic of the visual system, it does not explain the perception of contrast in complex images. This is because the CSF predicts visibility of very small, almost invisible contrast, presented on a uniform background, which is atypical for most complex scenes. The variations in contrast perception are much smaller for contrast sufficiently above the detection threshold. This was shown by Georgeson and Sullivan [1975], who measured the magnitude of contrast of one frequency that matches the magnitude of contrast of another frequency. They found that the lines of matching contrast across spatial frequencies range from a strongly bent curve for low contrast, which closely corresponds to the CSF, to an almost flat line for suprathreshold contrast. Georgeson and Sullivan coined the expression "contrast constancy" for the notion of the invariability of suprathreshold contrast across viewing conditions.

There is ample evidence that contrast constancy holds across the frequency range both for narrow-band patterns, such as sine-waves [Barten 1999] and for broadband patterns, such as bandpass-noise [Brady and Field 1995]. Brady and Field [1995] reported that contrast matches are almost perfect once the contrast is above the detection threshold without any gradual transition between near-threshold and suprathreshold vision. This, however, cannot be said about the contrast matches across the luminance range, where significant deviations from contrast constancy can be observed even for relatively large contrast magnitudes [Kulikowski 1976]. Therefore, we need to assume that the contrast constancy mechanism behaves differently in the frequency and luminance domains.

Kulikowski [1976] observed that, over a wide range of parameters, two contrast magnitudes match in their appearance when their supra-threshold contrast matches. That implies that the physical contrast M minus the detection threshold M_t must be equal for matching contrast:

$$M - M_t = \tilde{M} - \tilde{M}_t, \tag{4}$$

where M and $\tilde{M}$ are Michelson contrasts seen at different luminance. The detection threshold M_t is predicted by the CSF function:

$$M_t = \frac{\Delta L}{L} = \frac{1}{S \cdot CSF(\rho, L_a)}, \tag{5}$$

where ρ is the spatial frequency in cycles per degree and L_a is the adaptation luminance in cd/m^2. In our considerations we rely on the CSF from Mantiuk et al. [2011]. Also, similarly as in [Mantiuk

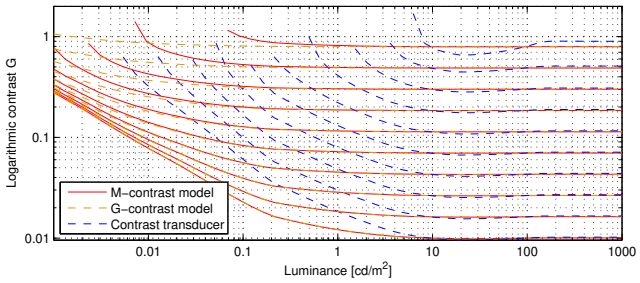

Figure 4: *The lines of matching contrast magnitude as a function of luminance. The lines join the contrast values that should appear the same according to the model.*

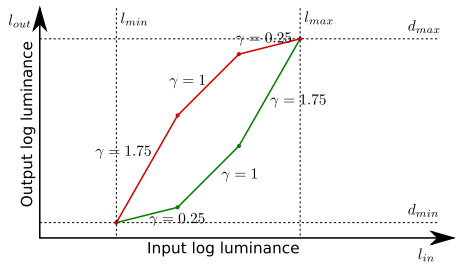

Figure 5: *Two piece-wise linear tone-curves. The green (lower) curve expands contrast in bright tones and compresses contrast in dark tones. Because the middle tones are pushed towards lower luminance levels, their perceived contrast is reduced. The opposite is achieved with the red (top) tone-curve. The slope γ describes the change of physical contrast.*

et al. 2011], we assume that an observer can locally adapt to the luminance of a pixel ($L_a = L$), which approximates the performance of the fast neural mechanism of adaptation. S is the absolute sensitivity factor, which is necessary to adjust the absolute thresholds for a particular experimental scenario. Using our experimental parameter adjustment setup (described in Section 3), we determined that $S = 8.6$ produces good matches. The peak sensitivity at 100 cd/m^2 for this S-value is $M_t = 0.4\%$, which is consistent with most CSF measurements [Barten 1999].

Although Kulikowski's model was defined in terms of Michelson contrast, it is convenient to formulate matching contrast in terms of the logarithmic contrast:

$$G - G_t = \widetilde{G} - \widetilde{G}_t, \quad \text{where} \quad G_t = G(M_t). \tag{6}$$

Note that Eq. 6 is not equivalent to Eq. 4 due to a non-linear relation between the contrast measures. However, as shown in Figure 4, matching contrast lines are almost identical for both models, except for very high contrast and low luminance. Because data does not exist for such high contrast levels, neither model can be said to be right or wrong. We will use the logarithmic contrast in our model because it does not suffer from singularities at high contrast.

Figure 4 also reveals an important characteristic of this contrast matching model. Lines of matching contrast are more curved for low contrast, which means that low contrast is more affected by luminance than high contrast. This is contrary to another popular model of suprathreshold contrast: contrast transducer [Pattanaik et al. 1998; Mantiuk et al. 2008]. The transducer predicts a much larger increase of physical contrast, regardless of the contrast magnitude. Such prediction is inconsistent with the experimental data.

Despite its simplicity, the model proposed by Kulikowski accurately predicts experimental data. In Figure 6 we collected contrast matching data from several sources and compared them with the model predictions. Kulikowski's model compares favorably to alternative models of perceived contrast, such as contrast transducer, models of brightness perception, JND luminance scaling, which all formed the lines very far from the data points (not included in the plot for better clarity). The model also encompasses our everyday experience of seeing in low light. The objects do not appear more blurry at night, as predicted by the multiplicative sensitivity loss in the aforementioned models. Instead their silhouettes are sharp but their textures lose low contrast details.

4.1.1 Global contrast

A tone curve is a powerful tool for reshaping image appearance. Its shape alters both physical and perceived image contrast, where the latter is affected by absolute luminance. To illustrate this, let us assume that any tone-curve can be approximated by a piece-wise linear function, such as the green (lower) curve shown in Figure 5. If we use the slope $\gamma = 1.75$ to expand contrast in brighter tones, we boost both perceived and physical contrast for these tones. But this also forces us to compress darker tones ($\gamma = 0.25$) as the dynamic range of the output device is limited to the range (d_{min}, d_{max}). Moreover, even though the physical contrast of the middle tones remains the same ($\gamma = 1$), the perceived contrast of these tones is lowered as they are pushed towards lower luminance. Therefore, in order to boost perceived image contrast, it is necessary to use an opposite tone-curve, such as one shown in red (top) in Figure 5. In this section, we demonstrate how to find a tone-curve that results in optimum perceived contrast given the limitations of the output device.

Our task is to find a tone-curve T that maps input luminance to output luminance so that the distortions in perceived contrast are minimized. We find such a curve for a representative contrast G and a spatial frequency ρ. For simplicity, the tone curve $T()$ is defined in the log luminance space

$$\widetilde{l} = T(l), \quad \text{where} \quad l = log_{10}(Y) \tag{7}$$

so that the resulting physical contrast can be expressed as:

$$\widetilde{G} = \frac{dT}{dl}\, G. \tag{8}$$

The above equation relies on the fact that the slope of a tone-curve in the log domain corresponds to the contrast change. The problem of finding the optimum tone-curve can be expressed as the optimization, where the squared difference of both sides of the Kulikowski's model (Eq. 6) is minimized. Formally, this can be expressed as:

$$\arg\min_{T(l)} \int_{l_{min}}^{l_{max}} S(l) \left(G - G_t(l) - \frac{dT}{dl}\, G + G_t(T(l)) \right)^2 + \\ + \tau \, (l - T(l))^2 \, dl \tag{9}$$

$$\text{subject to:} \quad \frac{dT}{dl} \geq 0, \tag{10}$$

$$T(l_{min}) \geq d_{min}, \quad T(l_{max}) \leq d_{max}. \tag{11}$$

$G_t(l)$ is the threshold contrast for log-luminance l (Eq. 5). The second term of the objective function is the difference between the source (l) and target log luminance ($T(l)$) and is weighted with a small constant $\tau = 0.0001$. The term is necessary to push the tone-curve towards either bright or dark tones when the dynamic

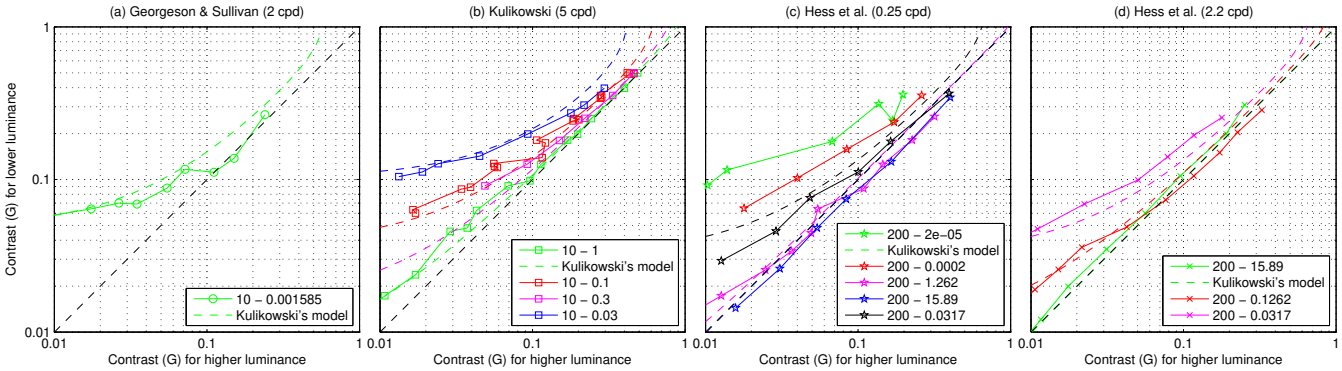

Figure 6: *Contrast matching data from several sources (continuous lines). Different colors denote different pairs of test and reference luminance (in cd/m^2) as listed in the legends. In general, a higher physical contrast is needed at lower luminance (y-axis) to match the appearance of contrast at higher luminance (x-axis). The predictions of Kulikowski's model, where available, are plotted as dashed lines with their color indicating the same pair of test and reference luminance as the contrast matching results.*

range of the target image is lower than the dynamic range of a display. The first constraint (Eq. 10) ensures that the tone-curve is monotonic and increasing. The two remaining constraints (Eq. 11) ensure that the tone-curve does not exceed the minimum and maximum luminance of the target display (d_{min}, d_{max}). Note that the dynamic range and black level of the display are the parameters of our method. Therefore, it is possible to adjust results for displays of varying contrast and seen under varying ambient illumination.

The optional saliency function $S(l)$ is the most beneficial for high dynamic range images, which may contain small areas that greatly expand the dynamic range but do not form a salient part of an image. In such a case, it is preferred to choose a tone curve that will foremost match the appearance of the areas that form a significant part of the image. This is achieved by assigning weights to different luminance levels during optimization. In the simplest case, the function is a histogram of the input image. However, this would assign high weights to large, uniform areas so it is beneficial to further weigh the histogram by a measure of contrast. The disadvantage of using the saliency function is that the tone curve changes depending on image content. Even with some form of temporal filtering, this can lead to temporal color inconsistencies [Eilertsen et al. 2013]. Therefore, for video processing and the display-referred scenarios, we set all saliency weights to 1.

The above optimization problem can be efficiently solved numerically after converting a tone curve into a discrete piecewise linear function. The quadratic terms of the objective function let us express the problem as quadratic programming with inequality constraints. Because the threshold function G_t introduces nonlinearity, the quadratic problem needs to be solved iteratively, where the threshold function is approximated with its first order Taylor expansion in each iteration. Because there are very few optimized variables (usually about 20-30), the solver is efficient . If no saliency function is used, the solution can be precomputed per pair of source (l_{min}, l_{max}) and destination (d_{min}, d_{max}) luminance ranges. For simplicity, we solve this problem for a single representative spatial frequency $\rho = 2$ cpd, which approximately corresponds to the peak sensitivity of the visual system for a range of luminance levels (refer to Figure 3-right), and for a representative contrast $G = 0.4$. These values were found to produce the best matching results using our experimental setup (Section 3).

Several tone-curves computed for different source and target luminance levels are shown in Figure 7. Note that, when retargeting from 100 to 1 cd/m^2, the tone-curve becomes less steep (gamma<1) for bright tones and more steep for dark tones. This behaviour is

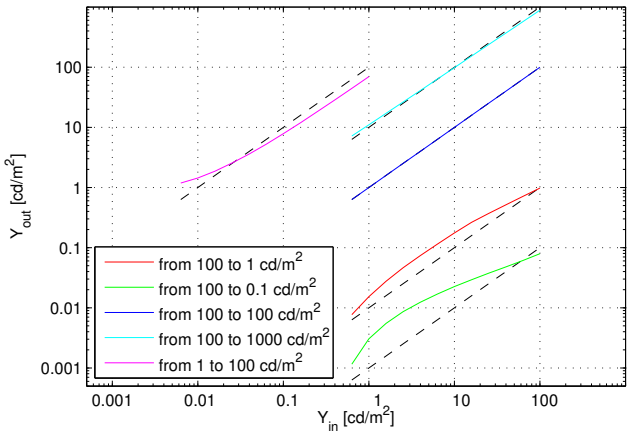

Figure 7: *The tone-curves for luminance retargeting that results in minimum perceived contrast distortion (refer to Section 4.1.1). The dashed lines represent linear mapping (gamma=1).*

very different from a typical gamma=1.5 curve used for "dark" conditions. There is also little change in the shape of the tone-curve when retargeting from 100 to 1000 cd/m^2, as the sensitivity (CSF) does not change much above 100 cd/m^2. The images produced by tone-curves that were optimized for different retargeting scenarios can be found in the top row of Figure 9.

Note that, as shown in Figure 2, the tone curve is applied to the full resolution luminance image in the *global contrast retargeting* step, followed by the extraction of a baseband. It may appear to be more efficient to apply the tone curve to the baseband extracted in the *laplacian pyramid decomposition* step. This, however, leads to strong halo artefacts when a nonlinear tone curve is applied to blurred edges in a base-band image.

4.1.2 Local contrast

A well selected tone curve can hugely improve the appearance of the retargeted image, however, it offers very coarse control over the contrast, limited to the selection of regions of similar luminance. Two other parameters of the contrast matching model, spatial frequency and contrast magnitude, must be addressed on a local level. To achieve local contrast control, we decompose an image into frequency-selective bands using the Laplacian pyramid (refer

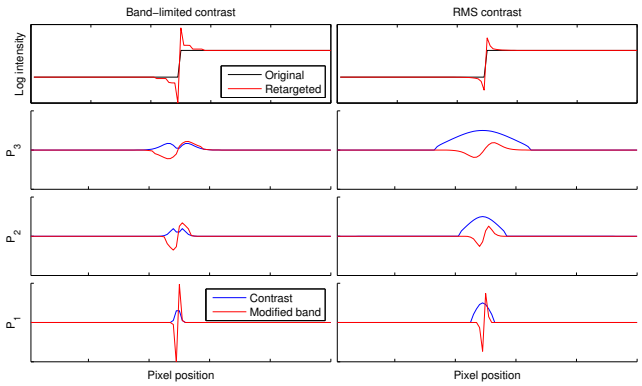

Figure 8: *An edge (black line, top) is enhanced with our local contrast retargeting method using either band-limited contrast (left) or RMS contrast (right). The plots labeled P_k show the band-pass or RMS contrast (blue) and the signal (red) in the band k after retargeting. Band-limited contrast underestimates the contrast of the edge and leads to excessive enhancement. RMS contrast can capture the contrast of an edge across the bands and does not cause over-enhancement.*

to Figure 2). The pyramid is computed for the log of luminance values so that the band-pass levels contain logarithmic contrast values (Eq. 2).

While spatial frequency is readily provided by the multi-scale decomposition, estimating contrast magnitude G requires more care. Contrast in complex images is typically estimated from the band-pass contrast representation [Peli 1990], which can be extracted from a Laplacian pyramid. However, there are two problems with this approach: a) Contrast is arguably best defined in terms of edges. Detecting edges, however, requires integration of information across several scales (frequency bands) [Witkin 1984]. Therefore, the perceived contrast is not formed by a single frequency band, but by integration of information from all frequency bands. b) Sharp edge contrast features are decomposed into smaller band-pass contrast components at several levels of the pyramid. These bandpass components are smaller than the total edge contrast and will be over-enhanced during retargeting to lower luminance level leading to errors in appearance mapping. This is visually illustrated in Figure 8.

We need a measure of contrast that integrates information from all frequencies, yet is localized and captures the contrast of a particular frequency band. A measure of contrast that is commonly used for noise and broad-band patterns is the root-mean-square (RMS) contrast:

$$c_{RMS} = \sqrt{\int \left(\frac{\Delta Y(\boldsymbol{x})}{\overline{Y}}\right)^2 d\boldsymbol{x}} = \sqrt{\int \left(\frac{Y(\boldsymbol{x}) - \overline{Y}}{\overline{Y}}\right)^2 d\boldsymbol{x}}, \tag{12}$$

where Y and ΔY are the image luminance and increment at the position $\boldsymbol{x}$, $\overline{Y}$ is the mean value, and the integral is computed over the entire image. The RMS contrast, however, gives a single value per stimulus and is not very useful for complex images. Therefore, we need to localize this measure by restricting it to a local Gaussian window. In order to relate the computed contrast measure to the logarithmic contrast, we operate on the log-luminance image $l = \log(Y)$ rather than luminance itself. Hence, the localized broadband contrast can be calculated as:

$$c(x,y) = \sqrt{\left(g_\sigma * [l(x,y) - (g_\sigma * l)(x,y)]^2\right)(x,y)}, \tag{13}$$

where $*$ is the convolution operator and g_σ is the Gaussian kernel with the standard deviation σ. The Gaussian window needs to get smaller for higher frequencies to account for finer scale. This is achieved by making it equal to half of the size of a single cycle at a particular frequency:

$$\sigma = 0.5\,\frac{R_{ppd}}{\rho}, \tag{14}$$

where R_{ppd} is the angular display resolution in pixels per visual degree and ρ is the spatial frequency in cycles per degree. σ is given in pixels assuming a non-decimated Laplacian pyramid, where all levels have the same resolution. The frequency ρ can be computed as:

$$\rho = 2^{-(k+1)}\,R_{ppd}, \tag{15}$$

where $k = 1,..,N$ is the level of the pyramid and $k = 1$ denotes the finest level.

Given the local contrast estimate, we perform contrast retargeting as a local enhancement of the Laplacian pyramid:

$$\widetilde{P}_k(x,y) = P_k(x,y) \cdot m_k(x,y) \tag{16}$$

where P_k corresponds to the source image pyramid level $k = 1..N$, and m_k is a contrast modification explained below. We select N so that the coarsest band (except the base band) has the peak frequency less or equal to 2 cpd. The low-pass base band ($k = N+1$) is discarded. The resulting image is reconstructed by summing all modified levels of the pyramid ($\widetilde{P}_k(x,y)$) including the base band, which comes from the global contrast retargeting step (refer to Figure 2). The contrast modification can be expressed as:

$$m_k(x,y) = \frac{c_k(x,y) - G(M_t) + G(\widetilde{M}_t)}{c_k(x,y)}, \tag{17}$$

where $c_k(x,y)$ is the contrast (Eq. 13) at the pixel location (x,y) and k-th level of the pyramid. The function $G()$ is contrast measure conversions, given in Eq. 3. M_t and $\widetilde{M}_t$ are the detection thresholds for the input and retargeted images (Eq. 5). In order to find these thresholds from the CSF, we use the peak frequency corresponding to the given band (Eq. 14) and pixel luminance of the source (Y) and retargeted ($\widetilde{Y}$) images. The latter is provided by the retargeted base-band image.

The result of the local contrast retargeting step isolated from other components of the algorithm can be seen in the second row of Figure 9. Note that the contrast is altered selectively depending on its magnitude. Such behaviour is consistent with the way we perceive contrast at different luminance levels.

4.2 Color retargeting

Reduced luminance affects not only luminance contrast but also color. This is manifested by loss of color saturation, mostly caused by reduced response of the cones, and the shift of hue towards more bluish colors, known as Purkinje shift. The latter effect is due to the fact that rods and cones share the same neural pathways to transmit their signal to the visual cortex [Cao et al. 2008]. In the photopic luminance range the information from the cones is the dominant signal, whereas in the scotopic range rods become dominant. In mesopic range, when both types of photoreceptor cells are active, the signal from the rods is combined with the signal from the cones in the early stages of visual processing. The variable contribution of the rod signal to the neural channel of each cone changes the ratio between the responses, resulting in the hue shift.

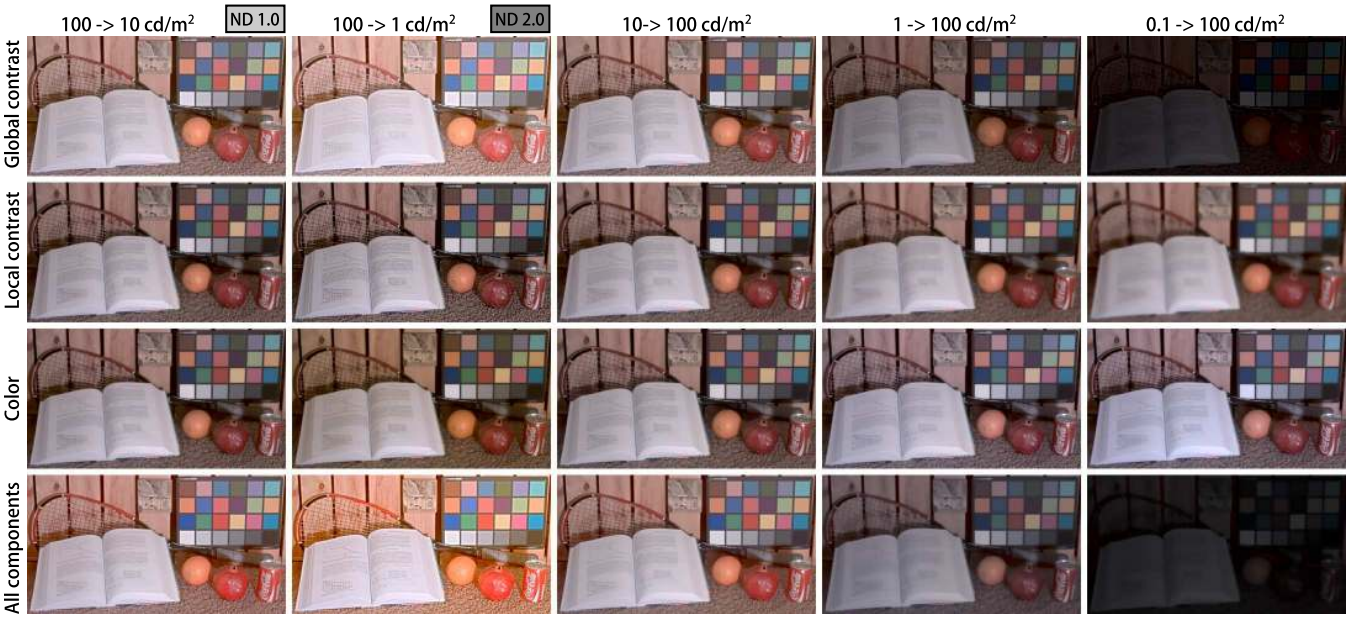

Figure 9: *Results produced by different components of our algorithm. The top row of numbers indicate source and target peak luminance of a display. Note that the results for retargeting to a dark display on the left (*100→10 *and* 100→1*) are meant to be seen at much lower luminance levels though a neutral-density filter as shown next to the label on top. Much of the apparent artefacts, such as haloing and over-sharpening, disappear when seen through an ND filter.*

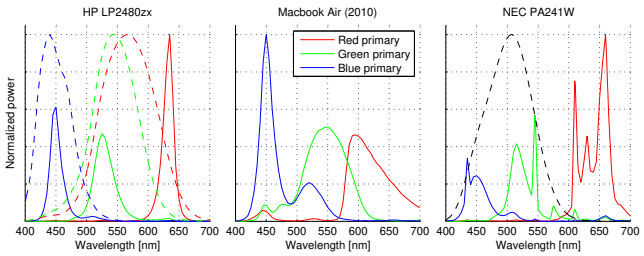

Figure 10: *Spectral emission of the tested displays. The left plot also shows Smith & Pokorny cone fundamentals (dashed lines), and the right plots shows the CIE scotopic luminous efficiency function (dashed black).*

Our goal is to find the resulting linear $[\widetilde{R}\ \widetilde{G}\ \widetilde{B}]'$ color values with a hue shift given input linear values $[R\ G\ B]'$ and the target luminance $\widetilde{Y}$. We start by modelling the response of the photoreceptor, which is the product of spectral distribution of light reaching the retina, $L(\lambda)$, and spectral sensitivity of each type of photoreceptor: L-, M-, S-cones and rods, $\sigma_P(\lambda)$:

$$E_P(C) = \int_\lambda L(\lambda)\,\sigma_P(\lambda)\,d\lambda, \tag{18}$$

where λ is the wavelength and index P corresponds to the type of photoreceptor: L, M, S, or R. We use the normalized Smith & Pokorny cone fundamentals [Smith and Pokorny 1975] for the L-, M- and S-cone sensitivities and CIE 1951 scotopic luminous efficiency function for rods. Usually, the incoming light is described as the product of three or more spectral basis functions (Π) and their coefficients (p):

$$L(\lambda) = \sum_{i=1}^{K} \Pi_i(\lambda)\,p_i. \tag{19}$$

Without losing generality, we can simplify the model and assume that the coefficients $p_{1..3}$ correspond to linear RGB color values. In Figure 10 we show the spectral primaries Π for several displays that we measured. It is then possible to find a matrix M_E for converting the linear RGB values into photoreceptor responses:

$$\begin{bmatrix} E_L \\ E_M \\ E_S \\ E_R \end{bmatrix} = M_E \begin{bmatrix} R \\ G \\ B \end{bmatrix}, \tag{20}$$

where the coefficients of the matrix M_E are given by:

$$m_{P,i} = \int_\lambda \Pi_i(\lambda)\,\sigma_P(\lambda)\,d\lambda. \tag{21}$$

Cao et al. [Cao et al. 2008] observed that the rod signal shares the pathway with L-, M-, and S-cone signals and its influence is additive and depends on the luminance of the signal. The combined responses of each cone channel with the rod input, L, M and S, can be expressed as:

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{pmatrix} 1 & 0 & 0 & k_1(Y) \\ 0 & 1 & 0 & k_1(Y) \\ 0 & 0 & 1 & k_2(Y) \end{pmatrix} \begin{bmatrix} E_L \\ E_M \\ E_S \\ E_R \end{bmatrix} = M_C(Y) \begin{bmatrix} E_L \\ E_M \\ E_S \\ E_R \end{bmatrix}, \tag{22}$$

where $k_1(Y)$ and $k_2(Y)$ are the functions modeling rod input strength to the L, M (k_1) and S (k_2) channels at luminance Y. These functions are obtained by linearly interpolating between the values measured in [Cao et al. 2008], which are listed in the table below.

Y [cd/m^2]	10	0.62	0.10
k_1	0	0.0173	0.173
k_2	0	0.0101	0.357

The signal is then processed further down the visual cortex and combined into opponent color space. However, since the transformation into the opponent color space is linear, we can match the

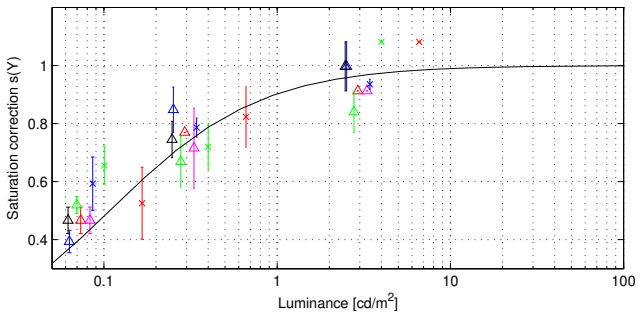

Figure 11: *Matching saturation factor with changing mean luminance of an image. The black line is the fitted curve (Eq. 26). Error bars denote standard errors.*

colors at this early stage. We assume that two colors at luminance Y and $\widetilde{Y}$ will appear similar if their relative cone response values are equal:

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} \cdot \frac{1}{Y} = \begin{bmatrix} \widetilde{L} \\ \widetilde{M} \\ \widetilde{S} \end{bmatrix} \cdot \frac{1}{\widetilde{Y}}, \tag{23}$$

Note that, while it is impossible to directly match LMS channels because of vastly different responses to a bright and dark display, the relative cone responses can be easily matched. After introducing Eq. 20 and 22 into Eq. 23, we can find the retargeted color values from:

$$\begin{bmatrix} \widetilde{R} \\ \widetilde{G} \\ \widetilde{B} \end{bmatrix} = \frac{\widetilde{Y}}{Y} (M_C(\widetilde{Y})\, M_E)^{-1} M_C(Y)\, M_E \begin{bmatrix} R \\ G \\ B \end{bmatrix}. \tag{24}$$

Matching cone contrast lets us correct for the hue shift but it does not account for the loss of color saturation caused by decreasing sensitivity of the cones as well as changes introduced by the tone-curve [Mantiuk et al. 2009]. We experimented with the complete model of [Cao et al. 2008], which introduces non-linear gains into opponent responses, but the results were unsatisfactory. The problem was that their model was generating colors that much exceeded the gamut of the target display. Instead, we found a simple saturation correction to work very well. After experimenting with saturation correction in CIE Lab, CIE Luv color spaces and a luminance-preserving method [Mantiuk et al. 2009], we found that the best results are produced by the common tone-mapping color correction formula:

$$\widehat{R} = \left(\frac{\widetilde{R}}{\widetilde{Y}} \right)^{\frac{s(Y)}{s(\widetilde{Y})}} \widetilde{Y}, \tag{25}$$

where R is the red trichromatic value. The same formula is applied to green and blue color channels. The matching saturation function $s(Y)$ was found in a matching experiment with a reference image shown at 200 cd/m^2 using the setup described in Section 3. The results of the experiment are shown in Figure 11 and the best fitted curve is given by:

$$s(Y) = Y/(Y + k_3), \tag{26}$$

where k_3 is equal to 0.108.

The result of color retargeting isolated from other components of the algorithm can be seen in the third row of Figure 9. Note that the hue changes due to Purkinje shift and loss of saturation at low luminance.

Figure 12: *Image compensated for viewing on a 10 cd/m^2 peak luminance display, individually for younger and older observers. Enlarge to see the difference.*

4.3 Timings

It takes approximately 5 seconds to process a single HD-resolution frame on an Intel Core i7 2.80GHz CPU using single-threaded Matlab code. This excludes the time of input/output operations and the time to calculate the tone curve, since it only has to be done once for each pair of displays and can be precomputed. Since most of the time is spent performing independent per-pixel operations, the algorithm is well suited for parallel processing.

5 Applications

Dark display. The primary application of our method is to compensate the appearance changes seen when images are shown on much darker displays. Figure 1 and 9 show examples of compensated images with higher overall brightness, boosted local contrast and increased color saturation. Such compensation is in particular attractive for mobile devices that can reduce their backlight illumination when used in dark environment thereby reducing power consumption. We found that the peak luminance of a 1000:1 display can be reduced to as little as 1 cd/m^2. Further brightness reduction results in excessive loss of color vision, which cannot be compensated. It is important to note that the compensation can take advantage of new display technologies, such as OLED, which offer much expanded color gamut and contrast. Such extra gamut can reproduce the highly saturated colors and contrast that can be found in compensated images.

Age-adaptive compensation. Because our method relies on a model of contrast sensitivity, it can be easily extended to account for the differences in acuity and sensitivity between young and elderly observers. In Figure 12 we show image compensation for a dimmed 10 cd/m^2 peak luminance display tailored for 20-year old and 80-year old observers. Typically little compensation is needed for 20-year old, but details and brightness must be boosted for the older observer.

Reproduction of night scenes. Our method can also retarget images of night scenes to reproduce their appearance on much brighter displays. Figure 9 shows examples of retargeting for a test scene and Figure 13 (center) shows an example for a scene-referred HDR image. Please note that the loss of acuity in the cathedral image is visible only in darker image parts, as expected. Although a number of tone-mapping operators and appearance models are meant to predict such appearance change, none of the existing methods can accurately predict the full range of effects, as discussed in the next section.

Creative rendering of night scenes. The actual appearance change due to low luminance is often subtle and much smaller than predicted by many visual models. To achieve more dramatic effect in entertainment applications, where perceptual accuracy is not crucial, it is often desirable to alter the appearance above the level

Figure 13: *Best exposure from a scene-referred HDR image (left) compared with a faithful simulation of night vision (center) and exaggerated rendering for a more dramatic effect (right). Compare the differences in the visibility of details and color.*

predicted by the visual model. This is shown in the right image of Figure 13, where we adjusted parameters to show an excessive change of image appearance.

Simulation of age-related vision loss. Similarly as it is possible to target dark-display compensation for an age-group, it is also possible to account for the age when reproducing night scenes. In Figure 14 we visualize a scene from a driving simulator, as seen by a 20- and 80-year-old observers. To complete simulation, we included in this application the age-dependent model of disability glare based on the CIE recommendation [Vos and van den Berg 1999].

Video. When content-independent approach is used ($S(l) = 1$ in Eq. 9), our method does not contain any temporarily inconsistent components and video can be processed frame-by-frame. A content-dependent approach requires temporal tone-curve filtering, such as the one proposed in [Mantiuk et al. 2008]. Examples of retargeted video clips can be found in the supplementary materials.

6 Comparison with other methods

In this section we compare our proposed method with several alternative techniques.

CIECAM02 is the state-of-the-art color appearance model, which accounts for a number of luminance-dependent effects, such as Hunt and Stevens effects (refer to Section 2). To retarget images, we process them through forward and then inverse CIECAM02 transforms. However, we vary the viewing-conditions-dependent parameters between the transforms. Depending on the source and target luminance levels, the viewing conditions vary between *dark*, *dim* and *average*. We also altered the luminance of the adapted white point to correspond to a drop in luminance levels, but we did not notice this parameter to have a significant effect on the results.

As shown in the top row of Figure 16, CIECAM02 predicts the loss of perceived contrast and color saturation at low light and compensates for it by boosting overall image contrast at the cost of reducing brightness ($100 \rightarrow 1$ cd/m^2 scenario). As we show later, such images offer an inferior appearance match. The appearance changes due to very low luminance ($1 \rightarrow 100$ cd/m^2 case) are too subtle, confirming that the model is in fact limited to the photopic vision.

Display adaptive tone-mapping [Mantiuk et al. 2008] is a tone-mapping operator that accounts for the dynamic range and absolute luminance of the target display. The operator utilizes a tone-curve optimization similar to retargeting global contrast in our algorithm, though based on the transducer function. The operator is limited to global (spatially invariant) tone-curve, which cannot account for frequency-dependent and color effects. We used the original implementation from the *pfstools/pfstmo* software.

Similarly as CIECAM02, the algorithm correctly predicts the loss of contrast with luminance (second row in Figure 16). However, it overpredicts the effect due to the transducer function. The colors that are too dark to be reproduced are clipped to black in the $100 \rightarrow 1$ scenario. The algorithm cannot retarget night scenes as it does not take into account the luminance of the input image.

Multi-scale model of adaptation, spatial vision and color appearance [Pattanaik et al. 1998] is one of the most comprehensive models of the visual system, accounting for a large range of appearance phenomena. We reimplemented the algorithm with the help of partial code fragments published by the authors. The best results were achieved when the low-pass band of the target image was multiplied by a constant factor, which is the treatment recommended by the authors for low-dynamic range images.

The results shown in Figure 16 demonstrate that the method predicts an extensive set of visual phenomena: loss of acuity, Purkinje color shift, loss of color saturation and contrast. However, it also clear that the magnitude of all these effects is not correctly predicted: the contrast and the acuity loss due to luminance is excessive, the color cast due to Punkinje shift is too subtle. The result for $100 \rightarrow 1$ reveal another limitation, shared with most forward-inverse visual models: the resulting colors often exceed the available dynamic range, resulting in a non-reproducible image.

Calibrated image appearance reproduction model [Reinhard et al. 2012] combines the goals of tone-mappings and color appearance to reproduce images on a range of display devices. We used the implementation provided by the authors and varied the input im-

Figure 14: *Simulation of night vision for 20 and 80 year old observers. The simulation assumes compensated refraction and the age-related vision loss due to reduced retinal illuminance (senile miosis and crystalline lens aging), disability glare, and neural sensitivity loss. Notice the loss of fine details (when enlarged on a screen), such as the car license number, in the image on the right. The driving simulator rendering is the courtesy of LEPSIS (part of IFSTTAR).*

Figure 15: *Comparison of our method with perceptual tone-mapping for low light conditions [Kirk and O'Brien 2011]. The image courtesy of Kirk and O'Brien. Enlarge to see the difference in detail visibility.*

age luminance and display adaptation according to the source and target luminance levels. The algorithm produces pleasing results over a wide variety of high dynamic range images. However, as shown in the 4^{th} row of Figure 16 there is little change in image appearance regardless of the retargeting scenario. This sugests that the model does not account for luminance-dependent effects in the non-photopic luminance range.

Perceptual mal-adaptation model [Irawan et al. 2005] is a tone-mapping operator that is capable of simulating loss of visibility experienced under changing illumination conditions. As shown in the fifth row in Figure 16, the method can predict reduced contrast and brightness for dark scenes. However, it does not contain spatial processing that could simulate loss of acuity, nor does it account for hue and saturation changes. The operator does not produce usable results when compensating for a dark display (100→1 scenario).

Tone-mapping for low-light conditions [Kirk and O'Brien 2011] employs the same model of Cao et al. [2008] as our method to simulate Purkinje shift. However, since the method assumes full adaptation to scotopic conditions across an image, it applies the same processing also to bright areas, which are seen by the photopic vision. The result is a bluish haze across the image shown in Fig. 15-center. Our method applies the hue shift selectively, only in the dark regions, producing images that more closely resemble the perception of night scenes. The method of Kirk et al. also does not simulate acuity loss, loss of cone sensitivity and the change of perceived contrast.

Our method is the most comprehensive model of luminance effects on vision from all the presented methods. It takes a very different strategy to global contrast retargeting and finds a tone-curve that obeys the constraints of the target display dynamic range, so that the resulting image does not suffer from excessive clipping of pixel values. The color cast due to Purkinje shift is visible, but only at low luminance levels. The local contrast modification does not simply sharpen or blur an image, but selectively reintroduces or removes image detail. The loss of acuity results in the loss of small contrast details while larger contrast is mostly unaffected. All these changes result in images that correspond to the actual image appearance when seen in our experimental set up, described in Section 3.

6.1 Experimental comparison

To objectively confirm that the proposed method offers a better appearance match, we ran a pairwise comparison experiment. From the methods discussed in the previous section, we selected only those that produced acceptable results in a particular retargeting scenario. We included a "gamma" function with the exponents 1.5, as this is common practice for dark viewing conditions [Fairchild 2005, p.125]. We also included the original unprocessed images as a control condition. The experimental setup was identical as the one described in Section 3, except that one portion of the screen contained two images, which were the result of two alternative retargeting methods. The observers were asked to choose the image that most closely matched the appearance of the image shown to the other eye when a 2.0D filter was worn on one eye or the other, depending on the scenario. 17 naïve observers, who did not take part in the parameter adjustment experiments, compared the methods for 8 scenes using the full pairwise design.

Results. In order to estimate what portion of the population selects one method as better than the other, the results were scaled in JND units using a similar method as in [Eilertsen et al. 2013]. The scaled results in Figure 17 show that our method was selected as providing a significantly better appearance match in almost all cases. Only in two cases, which were a portrait image *Woman* in the 200→2 scenario and *Flower* image in the 2→200 scenario, our method was comparable to the second best, though the ranking is not statistically significant. Surprisingly, very few of the existing methods provided reproduction better than the original unprocessed image. Even a contrast-enhancing gamma 1.5 seems to do more harm than good when retargeting for a dark display. Note, that we did not include the methods that did not work or failed in either retargeting scenario, such as a Display adaptive TMO in the 2→200 case and Mal-adaptation in the 200→2 case. These results clearly indicate that, unlike the existing algorithms, our method can produce consistently good results for two very different retargeting scenarios.

7 Limitations

Since our method relies on the contrast sensitivity function, it is as accurate as the CSF predictions. This limitation is shared with most perceptual methods based on threshold models. Individual observers differ in their contrast sensitivity, similarly as they differ in their visual acuity. If ones sensitivity differs from an average observer's sensitivity predicted by the CSF, the method may either under-compensate or over-compensate the image. If the source im-

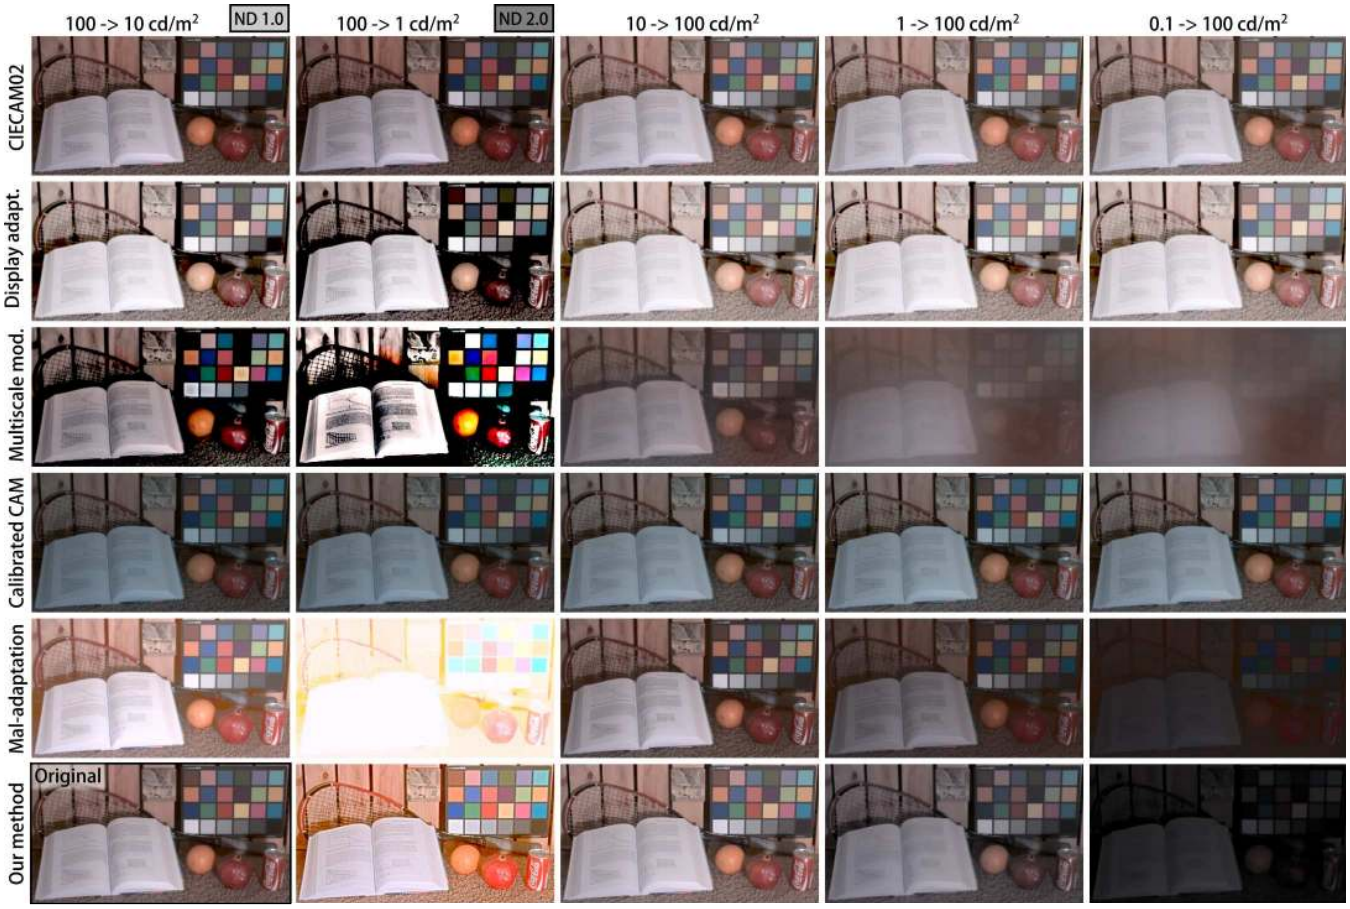

Figure 16: *The results of different methods (rows) when retargeting from one luminance level to another (columns). Columns labels are the same as in Figure 9. The original image is shown in the left-bottom corner instead of the* $100 \rightarrow 10$ *result for our method, which can be found in Figure 9. More results for a larger number of images can be found in the supplementary materials.*

age is noisy, the latter case can result in visible amplification of that noise. To prevent this, it may be desirable to adjust the CSF (parameter S in Eq. 6) for more conservative predictions.

Compensation for a very dark display may not produce an exact appearance match because the compensated colors are outside the available gamut. This problem can be partly addressed by using wide-color gamut displays and better gamut-mapping methods. We found that retargeting in the other direction, to simulate night vision, tended to provide more accurate results, as the image color gamut is shrunk and all colors are reproduceable.

Our model does not account for contrast and colorfulness changes above 200 cd/m^2, which are predicted by Stevens and Hunt effects. We attempted to measure these effects on an experimental HDR display using luminance levels up to 2 000 cd/m^2 but we did not find any measureable appearance changes.

8 Conclusions

Because human vision does not retain the same contrast and color perception across the luminance range, images need to be compensated when shown at a different luminance level than originally intended. Our method can provide such a compensation by retargeting night scenes for bright displays or retargeting bright scenes for a dark displays. The latter retargeting scenario allows for a novel application, in which an image is compensated for dark display, which leads to significant power saving in mobile devices while maintaining good image quality. Although many appearance models and tone-mapping operators claim to predict image appearance changes with luminance, we demonstrated that none of the existing models accounts for all relevant effects and can produce acceptable results for a range of luminance retargeting scenarios. While typical image appearance models usually involve a pair of forward and inverse perceptual models, which differ in the selection of viewing conditions, we take a very different approach with an optimized tone-curve. We bring from the field of vision science a simple but powerful supra-threshold contrast matching model, which has not been used in image and video applications before. The existing model of rod contribution to cone vision is used to predict Purkinje shift, and combined with our new measurements to predict also color saturation loss. Each component and the entire method is tested in experimental conditions to ensure a good appearance match.

Acknowledgements

We would like to thank Karol Myszkowski and Nigel John for their valuable comments. This work has been carried out within the Wales Research Institute of Visual Computing (RIVIC).

References

BARTEN, P. G. J. 1999. *Contrast sensitivity of the human eye and its effects on image quality*. SPIE Press.

BARTLESON, C. J., AND BRENEMAN, E. J. 1967. Brightness Perception in Complex Fields. *Journal of the Optical Society of*

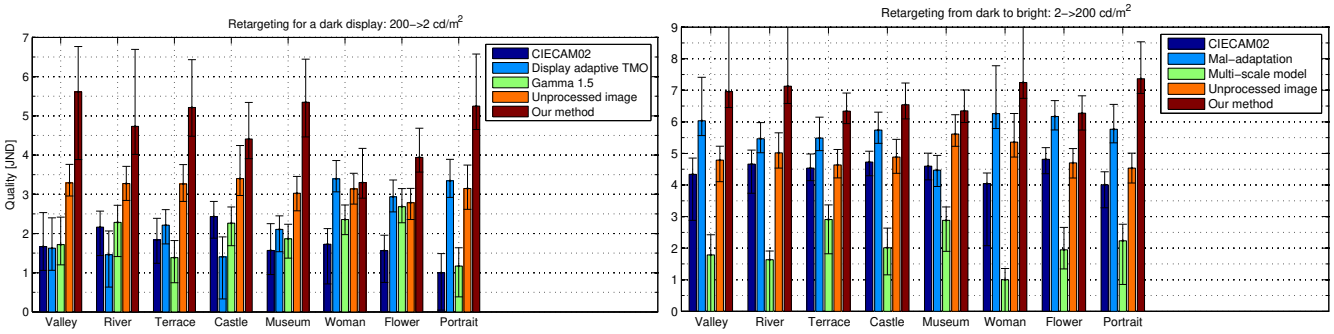

Figure 17: *The results of the pairwise comparison experiment scaled in JND units (the higher, the better) under Thurstone Case V assumptions, where 1 JND corresponds to 75% discrimination threshold. Note that absolute JND values are arbitrary and only relative differences are meaningful. The error bars denote 95% confidence intervals computed by bootstrapping. The images can be found in the supplementary materials.*

America 57, 7, 953–956.

BERNS, R. S. 1996. Methods for characterizing crt displays. *Displays 16*, 4, 173–182.

BRADY, N., AND FIELD, D. J. 1995. What's constant in contrast constancy? The effects of scaling on the perceived contrast of bandpass patterns. *Vision Research 35*, 6, 739–756.

CAO, D., POKORNY, J., SMITH, V. C., AND ZELE, A. J. 2008. Rod contributions to color perception: linear with rod contrast. *Vision Research 48*, 26, 2586–92.

CHANG, N., CHOI, I., AND SHIM, H. 2004. DLS: Dynamic backlight luminance scaling of liquid crystal display. *IEEE Trans. on Very Large Scale Integration Systems 12*, 8, 837–846.

EILERTSEN, G., WANAT, R., MANTIUK, R. K., AND UNGER, J. 2013. Evaluation of Tone Mapping Operators for HDR-Video. *Computer Graphics Forum 32*, 7 (Oct.), 275–284.

FAIRCHILD, M. D. 2005. *Color Appearance Models*, 2nd ed. John Wiley & Sons.

GEORGESON, B. Y. M. A., AND SULLIVAN, G. D. 1975. Contrast constancy: deblurring in human vision by spatial frequency channels. *The Journal of Physiology 252*, 627–656.

IRANLI, A., LEE, W., AND PEDRAM, M. 2006. Backlight dimming in power-aware mobile displays. In *Proc. of the Annual Conference on Design Automation - DAC '06*, 604–607.

IRAWAN, P., FERWERDA, J., AND MARSCHNER, S. 2005. Perceptually based tone mapping of high dynamic range image streams. In *Proc. of EGSR*, 231–242.

KANE, J. P., AND KURTZ, F. A., 2012. Adapting display color for low luminance conditions. Patent app. EP 2526688 A1.

KEROFSKY, L., AND DALY, S. 2006. Brightness preservation for LCD backlight dimming. *Journal of the Society for Information Display 14*, 12, 1111–1118.

KIRK, A. G., AND O'BRIEN, J. F. 2011. Perceptually based tone mapping for low-light conditions. *ACM Trans Graph 30*, 4, 42:1–42:10.

KUANG, J., JOHNSON, G., AND FAIRCHILD, M. 2007. iCAM06: A refined image appearance model for HDR image rendering. *Journal of Visual Communication and Image Representation 18*, 5, 406–414.

KULIKOWSKI, J. J. 1976. Effective contrast constancy and linearity of contrast sensation. *Vision Research 16*, 12, 1419–31.

MANTIUK, R., DALY, S., AND KEROFSKY, L. 2008. *ACM Trans Graph 27*, 3, 68:1–68:10.

MANTIUK, R., MANTIUK, R. K., TOMASZEWSKA, A., AND HEIDRICH, W. 2009. Color correction for tone mapping. *Computer Graphics Forum 28*, 2, 193–202.

MANTIUK, R., KIM, K. J., REMPEL, A. G., AND HEIDRICH, W. 2011. HDR-VDP-2: A calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Trans Graph 30*, 4, 40:1–40:14.

MORONEY, N., FAIRCHILD, M., HUNT, R., LI, C., LUO, M., AND NEWMAN, T. 2002. The CIECAM02 color appearance model. In *Proc. IS&T/SID 10th Color Imaging Conference*, Society for Imaging Science and Technology, 23–27.

PATTANAIK, S. N., FERWERDA, J. A., FAIRCHILD, M. D., AND GREENBERG, D. P. 1998. A multiscale model of adaptation and spatial vision for realistic image display. In *Proc. of SIGGRAPH '98*, 287–298.

PELI, E. 1990. Contrast in complex images. *Journal of the Optical Society of America A 7*, 10, 2032–2040.

PETIT, J., AND MANTIUK, R. K. 2013. Assessment of video tone-mapping : Are cameras S-shaped tone-curves good enough? *Journal of Visual Communication and Image Representation 24*, 1020–1030.

REINHARD, E., POULI, T., KUNKEL, T., LONG, B., BALLESTAD, A., AND DAMBERG, G. 2012. Calibrated image appearance reproduction. *ACM Trans Graph 31*, 6, 201:1–201:11.

SMITH, V. C., AND POKORNY, J. 1975. Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm. *Vision Research 15*, 2, 161–171.

THOMPSON, W., SHIRLEY, P., AND FERWERDA, J. 2002. A spatial post-processing algorithm for images of night scenes. *Journal of Graphics Tools 7*, 1, 1–12.

VOS, J. J., AND VAN DEN BERG, T. J. 1999. CIE 135/1-6 Disability Glare. Tech. rep.

WITKIN, A. 1984. Scale-space filtering: A new approach to multi-scale description. In *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing*, vol. 9, 150–153.