

Apparent Greyscale: A Simple and Fast Conversion to Perceptually Accurate Images and Video

Kaleigh Smith, Pierre-Edouard Landes, Joëlle Thollot, Karol Myszkowski

► **To cite this version:**

Kaleigh Smith, Pierre-Edouard Landes, Joëlle Thollot, Karol Myszkowski. Apparent Greyscale: A Simple and Fast Conversion to Perceptually Accurate Images and Video. Computer Graphics Forum, Wiley, 2008, Special Issue: Proceedings of Eurographics 2008, 27 (2), pp.193-200. 10.1111/j.1467-8659.2008.01116.x . inria-00255958

HAL Id: inria-00255958

<https://hal.inria.fr/inria-00255958>

Submitted on 14 Feb 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Apparent Greyscale: A Simple and Fast Conversion to Perceptually Accurate Images and Video

Kaleigh Smith¹ Pierre-Edouard Landes² Joëlle Thollot² Karol Myszkowski¹

¹MPI Informatik, Saarbrücken, Germany ²Grenoble Universities, INRIA, France

Abstract

This paper presents a quick and simple method for converting complex images and video to perceptually accurate greyscale versions. We use a two-step approach first to globally assign grey values and determine colour ordering, then second, to locally enhance the greyscale to reproduce the original contrast. Our global mapping is image independent and incorporates the Helmholtz-Kohlrausch colour appearance effect for predicting differences between isoluminant colours. Our multiscale local contrast enhancement reintroduces lost discontinuities only in regions that insufficiently represent original chromatic contrast. All operations are restricted so that they preserve the overall image appearance, lightness range and differences, colour ordering, and spatial details, resulting in perceptually accurate achromatic reproductions of the colour original.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Display Algorithms

1. Introduction

The basic problem of greyscale transformation is to reproduce the intent of the colour original, its contrasts and salient features, while preserving the perceived magnitude and direction of its gradients. The transformation consists of two interdependent tasks: a mapping that assigns a grey value to each pixel or colour, and a discriminability constraint so that the achromatic differences match their corresponding original colour differences. Recent approaches solve discriminability constraints to determine the grey values, producing images in which the original colour contrasts are highly discriminable. However, the greyscale images may exhibit exaggerated dynamic range, an arbitrary achromatic order that differs among colour palettes, and a smoothing or masking of details. These modifications all contribute to make the grey version appear dissimilar from its original and create inconsistency among like images and video frames.

The goal of this work is to create a perceptually accurate version of the colour image that represents its psychophysical effect on a viewer. Such greyscale imagery is important for printed textbooks and catalogues, the stylization of videos and for display on monochromatic medical displays. A perceptually accurate image is one that emulates both global and local impressions: it matches the original values' range and average luminance, its local contrasts are neither ex-

aggerated nor understated, its grey values are ordered according to colour appearance and differences in spatial details are imperceptible. Strong perceptual similarity is particularly important for consistency over varying palettes and temporal coherence for animations.

We present a new two-step greyscale transformation that combines a global mapping based on perceived lightness with a local chromatic contrast enhancement. Our simple algorithm yields comparable images to more complex approaches, and its linear runtime makes it suited to video processing and accelerated graphics hardware implementations. First, the grey values are mapped pixelwise from each colour's apparent lightness, resulting in the reproduction of the original contrast and gradients. Second, the gradient magnitudes are measured in perceptual difference ΔE and adjusted to maintain or improve discriminability with our multiscale chromatic contrast enhancement filter. This twofold approach mimics aspects of the human visual system, which processes global attributes while simultaneously depending on local contrasts such as edges and surrounds.

We choose *lightness* as the quantity for the grey values (and thus their ordering) because it is the achromatic response to a colour stimulus, measuring how bright a colour appears compared to an equally bright white. Colour studies show that lightness depends largely on luminance, but that colour-

fulness also contributes, as characterized by the Helmholtz-Kohlrausch effect (H-K); a colourful stimulus appears more light than a similar less colourful sample. The H-K effect has been identified as an important factor in greyscale mapping, and although it has been used for clipart greyscale mapping [BB04], no existing greyscale conversion for complex images explicitly takes it into account. Our *global apparent lightness mapping* is independent of the original colour palette, incorporates the H-K effect so is sensitive to small differences even between isoluminant colours, and yields perceptually accurate gradient directions and an appropriate dynamic range.

The mapping from a 3D to 1D colour space reduces the overall difference between colours, jeopardizing discriminability. We are not too sensitive to this loss when it occurs between spatially distant colours, but with adjacent colours it is immediately apparent, especially if an original contrast becomes imperceptible. To solve this problem, we enhance local contrast until its magnitude emulates that in the original. The enhancement restores chromatic differences without overemphasizing luminance differences by adaptively increasing weak contrasts. Furthermore, the process is restricted so that the polarity over edges, overall lightness and colour ordering are preserved, thus maintaining perceptual accuracy.

We begin by describing other greyscale conversion techniques and highlighting their effects on perceptual accuracy. Then, in Section 3, we give a background on colour appearance models, compare H-K predictors and determine the best suited to greyscale conversion. In Sections 4 and 5, we present our global-local technique to solve greyscale conversion. In Section 6, we compare the results of our technique to other standard results to accentuate the perceptual aspects that are more accurately preserved, and demonstrate our technique's ability on various input types. Finally, we conclude by discussing the implications of our findings, future work and other possible applications.

2. Related Work

There are a variety of printing and display solutions catered to the conversion of images from colour to greyscale. The most straightforward conversion maps a colour to an equi-luminant grey value, by desaturation or by picking a single colour channel to mimic the effect of a colour filter.

In their short paper studying chromatic contrast for greyscale conversion, Bala et al. [BE04] take a spatial approach and introduce color contrasts in CIELAB *LCH* (lightness, chroma, hue angle) by adding the high-pass filtered chroma channel to the lightness channel. To prevent overshooting in already bright areas, this correction signal is locally adjusted and its sign is taken from the lightness contrast. The algorithm is susceptible to problems in chroma and lightness misalignment. Taking a local adaptive approach, Bala et

al. [BB04] propose a mapping method for business graphics. The distinct colours of the image are sorted according to a simplified lightness predictor that incorporates the H-K effect. To maximize discriminability, adjacent pairs of lightness values are then respaced according to their relative color differences. The approach is uniquely for graphics with up to 10 colours, and is not applicable to complex images.

Gooch et al. [GOTG05] find grey values that best match the original color differences through an objective function minimization process. Original contrast between each pixel and its neighbours is measured by a signed distance, whose magnitude accounts for luminance and chroma difference and whose sign represents the hue shift with respect to a user-defined hue angle. It has $O(N^2)$ to $O(N^4)$ complexity, but a recent extension to a multiresolution framework by Mantiuk et al. improves the algorithm's performance [MMS06]. Rasche et al. [RGW05] propose a similar approach that finds the linear transform matching pairwise grey differences to corresponding color differences. The best transform is found by minimizing an error function that can be evaluated over a smaller set of colors to alleviate computation costs.

In recent work, Grundland et al. [GD07] find a global continuous mapping that adds lost chromatic information to the luminance channel. Their algorithm achieves linear-time performance thanks to Gaussian pairing sampling which limits the amount of processed color differences. In *YPQ* color space, the color differences are projected onto the two predominant chromatic contrast axes and are then added to the luminance image. A saturation-controlled adjustment of the output dynamic range is adaptively performed to balance between the original range and the desired amount of enhancement. Recently, Neumann et al. [NCN07] present a technique with linear complexity that requires no user intervention. It stresses perceptual loyalty by measuring the image's gradient field by colour differences in their *Coloroid* color space. After discarding all gradient field inconsistencies, fast 2D integration determines the final greyscale image.

In [CF03], Calabria and Fairchild find that image lightness strongly affects perceived contrast, meaning techniques that can arbitrarily modify lightness, like approaches by Rasche and Grundland, may affect image appearance in an adverse way. A greyscale ordering that contradicts the colours' luminance ordering also strongly impacts image appearance, yet in several approaches, ordering is subjective and arbitrary: the choice of hue angle in Gooch's Color2Gray can change all gradient directions, in Rasche's approach a user-defined threshold controls whether a colour is mapped to a darker or lighter value (see Figure 7), and in Grundland's approach ordering depends on the image and parameter choice of the color sampling method. Lastly, image details and salient features may be lost by the choice of neighbourhood size in Gooch's Color2Gray or by unpredictable behavior in inconsistent regions of the gradient field in Neumann's approach

(see Figure 9). Since the discussed colour to grey methods depend strongly on local image content, colour palettes and user parameters, they are hindered in their perceptual accuracy and are not directly applicable to animation where the colour palette is frequently modified and pixel correlations change quickly due to occlusion and disocclusion. With this as motivation, we attempt a different approach where perceptual accuracy and consistency is paramount.

3. Apparent Lightness

In order to discuss colour, we must move into the realm where colour exists - that is, in the observer's mind. A physical stimulus produces a perceptual response that we name 'colour', or alternately, an achromatic response we name 'brightness' or 'perceived lightness'. Colour appearance models take on the complex task of predicting a human viewer's perceptual response to colour stimulus - thus defining measures of colour. Throughout this paper, we work in the *CIELAB* and *CIELUV* colour spaces, whose three axes approximate perceived lightness L^* , chroma C^* and hue angle H^* . The first component, L^* , quantifies the perceptual response of a human viewer to luminance and is defined by Hunt as "the brightness of objects relative to that of a similarly illuminated white". Mathematically, it is defined as $L^* = 116(Y/Y_0)^{1/3} - 16$ for luminance Y and reference white luminance Y_0 .

While luminance is the dominant contributor to lightness perception, the chromatic component also contributes, and this contribution varies according to both hue and luminance. For example, cornflower blue seems brighter than a dull golden yellow of equal luminance. This phenomenon is characterized by the Helmholtz-Kohlrausch effect, where given two isoluminant colours, the more colourful sample appears brighter.

Helmholtz-Kohlrausch Effect A chromatic stimulus with the same luminance as a white reference stimulus will appear brighter than the reference [Nay97].

There are two experimental approaches for measuring the H-K effect: the Variable-Achromatic-Colour (VAC) approach, in which an achromatic sample's luminance is adjusted to match a colour stimulus; and the Variable-Chromatic-Colour (VCC) approach, in which the chromatic content of a colour stimulus is adjusted until its brightness matches a given grey stimulus [Nay98]. VAC is more common and was used in the seminal 1954 Sanders-Wyszecki study, and again in Wyszecki's later 1964 and 1967 studies [Wys67].

3.1. Helmholtz-Kohlrausch Lightness Predictors

The H-K phenomenon is predicted by a *chromatic lightness term* that corrects L^* based on the colour's chromatic component. We examine three such predictors published by

Fairchild and Nayatani for suitability to the greyscale problem. Existing models, like CIECAM02, account for many more complex colour appearance aspects, like surrounding colours, but are less suited to greyscale conversion due to their complexity and because most disregard the Helmholtz-Kohlrausch effect (the reader may refer to Table 17.1 in [Fai05]).

Fairchild's *CIELAB* chromatic lightness metric L^{**} is fit to Wyszecki 1967 data and is defined as [FP91]:

$$L^{**} = L^* + (2.5 - 0.025L^*) \left(0.116 \left| \sin \left(\frac{H^* - 90}{2} \right) \right| + 0.085 \right) C^* \quad (1)$$

Chroma C^* measures colourfulness and a sinusoidal curve predicts the H-K effect's decreased impact at yellow hues and its strongest effect at blues.

Nayatani defines chromatic lightness metrics $L_{N_{VAC}}^*$ and $L_{N_{VCC}}^*$, for each experimental approach, based in *CIELUV* [Nay97]. † A quantitative difference between them is that $L_{N_{VCC}}^*$ is twice as strong $L_{N_{VAC}}^*$ (in log space). For each method, chromatic object lightness is predicted by the following equations (see Appendix for more details):

$$L_{N_{VAC}}^* = L^* + [-0.1340 q(\theta) + 0.0872 K_{Br}] s_{uv} L^* \quad (2)$$

$$L_{N_{VCC}}^* = L^* + [-0.8660 q(\theta) + 0.0872 K_{Br}] s_{uv} L^* \quad (3)$$

s_{uv} is the chromatic saturation in terms of u', v' which predicts the effect's strength according to colourfulness. The quadrant metric $q(\theta)$ predicts the change of the H-K effect for varying hues, and constant K_{Br} expresses the H-K effect's dependance on the adapting luminance L_a . These chromatic




Figure 1: Lightness values from various H-K effect predictors applied to a spectrum of isoluminant colours, compared to CIE L^* .

lightness metrics solve a key challenge in greyscale conversion because they predict differences between isoluminant

† For readability, we have used the notation from [Nay98].

colours. Figure 1 plots the lightness measured by each metric on a nearly equiluminant colour ramp. It can be seen that more variation occurs when the H-K is being predicted, compared to luminance-based L^* which predicts nearly equal lightness for all colours. Note that other colour pairs will map to the same greyscale value, but that these are predicted to be more similar than the isoluminant colours.

We now decide which predictor is best suited to greyscale conversion. We prefer L^{**} or L_{NVAC}^* , because the gathering of VAC data on which they are modeled seems more akin to the goal as it finds a grey that matches a colour. Moreover, in testing L_{NVCC}^* , we observe that its stronger effect maps many bright colours to white, making it impossible to distinguish between very bright isoluminant colours [Nay98]. For that reason, and by heeding Nayatani's advice that L_{NVAC}^* , instead of L_{NVCC}^* , should be used for predicting differences between isoluminant colours, we decide not to use L_{NVCC}^* [Nay98].




Figure 2: Lightness values from various H-K effect predictors applied over a full spectrum. L^{**} exhibits a small range and at blue hues differs from L^* .

Because they are both fit to VAC data, the behaviours of L^{**} and L_{NVAC}^* are very similar. Their differences stem from the data on which they are based, and the flexibility of the models. L_{NVAC}^* is based on both Wyszecki 1964 and 1967 data, theoretical arguments about H-K, and the effect of adapting luminance. The L^{**} model is based only on Wyszecki 1967 data and has a simpler treatment of hue which we expect is responsible for the following: we observed L^{**} of blue hues is much higher than L^* , which reduces its range and makes its ordering differ significantly from both L_{NVAC}^* and L^* , see Figure 2. While the model fits the H-K effect perceptual data, this range reduction is problematic for greyscale conversion because colours with different L^* become less discriminable, an observation shared by Bala [BB04].[‡] We therefore conclude that L_{NVAC}^* is the most suitable H-K predictor to use in our global colour to greyscale mapping.

[‡] Bala uses $L_1^{**} = L^* + 0.143C^*$.

4. Global Apparent Lightness Mapping

We now describe our global mapping according to apparent lightness using the Nayatani model $L_N^* = L_{NVAC}^*$ described in the previous section. The mapping process is as follows:

$$I_{RGB} \rightarrow I_{LUV} \rightarrow I_{L_N^*} \rightarrow G \quad (4)$$

We first convert the colour image to linear RGB by inverse gamma mapping, then transform to *CIELUV* colour space. Its apparent chromatic object lightness channel L_N^* is calculated according to Equation 2. We map L_N^* to greyscale Y values using reference white chromatic values for u^* and v^* . Finally, we apply gamma mapping to move from linear Y space back to a gamma-corrected greyscale image G . As shown in Figure 3 for several colour ramps, the mapping is continuous, there is no colour reordering, no lost discrimination and the dynamic range is preserved.




Figure 3: On a colour test (left), G (right) preserves overall appearance and lightness ordering.




Figure 4: Our approach maps isoluminant colours to unique, properly ordered greyscale values.

Due to the compression of a 3D gamut to 1D, L_N^* may map two different colours to a similar lightness, which then are quantized to the same grey value. This occurs only when colours differ uniquely by hue, which is very uncommon in natural images and well-designed graphics. Even for a very challenging image that comprises equiluminant colours sampled from Neumann et. al's paper [NCN07], our global mapping discriminates appropriately, predicting the H-K effect that makes a more colourful blue appear lighter than the duller yellow, as shown in Figure 4 (view original colours on a calibrated screen). Recall that our goal is perceptual accuracy: the resulting low contrast properly represents the low contrast of the colour image, and each unique colour is mapped to a unique greyvalue. By incorporating the H-K effect, our global mapping partially solves the problem of grey value assignment and appropriately orders colours that normal luminance mapping can not discriminate.

5. Local Chromatic Contrast Adjustment

The mapping described in the previous section captures some chromatic content in the greyscale image G according to the H-K effect. However, because of dimension reduction and unaccounted for hue differences, chromatic contrast may be reduced. Humans are most sensitive to these losses at local contrasts, regions where there is a visible discontinuity. To counter the reduction, we increase local contrast in the greyscale image G to better represent the local contrast of original I . The technique is adapted according to the ratio between colour and greyscale contrast, so that increases occur at underrepresented colour edges without unnecessarily enhancing edges that already represent the original. We cater the general adaptively-weighted multiscale unsharp masking technique [NB98] to our goal of reintroducing chromatic contrast. Recently, this general tool has been used to good effect for adjusting contrast of tone mapped images [KMS07] and to restore contrast from depth perception [LCD06].

We perform contrast adjustments using the Laplacian pyramid that decomposes an image into n bandpass images h_i and a single lowpass image l [BA83]. Laplacian pyramids are built for I and G in CIE_{LAB} using a binomial coefficient filter of order 4. The h_i of each channel measures its local contrast, but as G contains no chromatic information, its local contrasts are contained entirely in its L^* channel. At each scale in the Laplacian pyramid, we adaptively increase local contrast $h_i(G_{L^*})$ by a perceptually-based amount λ_i , which measures the amount of contrast needed to match colour contrast $h_i(I)$. The enhanced greyscale image G' is computed by modifying G_{L^*} as follows:

$$G'_{L^*} = G_{L^*} + \sum_{i=0}^{n-1} k_i \lambda_i h_i(G_{L^*}) \quad (5)$$

where parameters k_1, \dots, k_{n-1} , $k_i \leq 1$ exist so that the spatial effect can be controlled according to amount of discriminability desired and the intended viewing conditions (image size and viewing distance).

The goal of gain factor λ_i is to measure the remaining chromatic contrast to be restored during the enhancement. We define it as:

$$\lambda_i = \left(\frac{\Delta E(h_i(I))}{|h_i(G_{L^*})|} \right)^p \quad (6)$$

$\Delta E(h_i(I))$ is the colour contrast between a pixel and its neighbourhood which we measure by $\Delta E(h_i(I)) = (h_i(I_{L^*})^2 + h_i(I_{a^*})^2 + h_i(I_{b^*})^2)^{\frac{1}{2}}$. Since the chromatic channels of G contain no contrast information, $|h_i(G_{L^*})| \cong \Delta E(h_i(G))$. The ΔE colour difference is used so that both colour and grey contrasts are expressed in units of perceptual lightness. The parameter $0 \leq p \leq 1$ is used to remap the λ values to a non-linear scale so that weaker contrasts, like those from isoluminant colours, can be enhanced without over emphasizing stronger contrasts.

The parameters introduced exist to provide flexibility, allowing users to tweak according to their preference for desired discriminability. We criticize other approaches for being ad hoc, so have ensured that the parameters provide flexibility without allowing uncontrolled changes to the image. The overall lightness is not altered because we limit the number of subbands that may be enhanced, preventing changes from having too large an impact (in practice $n \leq 4$ levels). Most importantly, by definition of λ , edge polarity can not flip, meaning the lightness order of adjacent regions can not be changed.




Figure 5: Compared to the basic unsharp mask (middle), our chromatic enhancement (right) gains contrast where it is low in G and high in I . In the gain images, green values represent negative gain; $p = 0.25$ $k = \{2, 1, 1, 0.6\}$.

We illustrate the effect of our local chromatic contrast adjustment in Figure 5. Contrast is nearly below threshold between isoluminant regions, especially among the bottom row of colours. A basic sharpening of all contrast (middle) does little to discriminate along that bottom row because the bandpass contains next to no signal. With our chromatic adjustment (right) it is possible to lift these contrasts above threshold without over emphasizing existing contrasts, so the resulting image better represents the original contrast.

6. Results and Discussion

The results presented here strive for perceptual accuracy, and do not attempt to increase or exaggerate discriminability. Therefore, the effects are apparent, but subtle. For comparison, we present either the CIE Y channel or Gimp greyscale with a basic unsharp enhancement, so that the reader is able to compare between images with matching overall sharpness. Additionally, the images presented here are for viewing on a calibrated colour screen (sRGB); for print, our resulting greyscale images should be mapped to the appropriate printer gamut.

We begin by showing that we discriminate between isoluminant colours by applying our approach to two images from Gooch et al. [GOTG05] in Figure 6. In both images, nearly

isoluminant regions are distinguishable (parameters p and k , and thus n , are given with each image). We now compare to previous work to support our claims of better perceptual accuracy. We illustrate the consistency problem of local adaptive approaches in Figure 7 with Rasche’s result using default parameters (top row) on similar images with varying colours (bottom row). While our results (middle row) also use default parameters $p = 0.5, k = \{0.5, 0.5, 0, 0\}$, the correct flower brightness ordering is preserved and the grey-values of leaves and backgrounds are identical. Additionally, our images do not exaggerate the overall lightness range. Because we maintain consistency, we are able to apply our approach to video, as shown for a single frame in Figure 8 with constant parameters $p = 0.8, k = \{0.2, 0.8, 0, 0\}$. The red flowers become more visible and bright without changing the overall video appearance and maintaining temporal coherence. Please see the accompanying video or our project website for the full animation and other examples of our approach applied over time [WWW08].




Figure 7: Consistency and colour ordering: Rasche local conversion top, our conversion (middle), original colour (bottom) are recoloured versions from Rasche.




Figure 6: Our conversions discriminate between colours even with similar luminance. Car $p = 0.5, k = \{0.5, 0.5, 0.5\}$; Island $p = 0.8, k = \{0.4, 0\}$.

We consider changes to spatial content, as illustrated by greyscale versions of Monet’s Impression Sunrise shown in Figure 9. Gooch’s Color2Gray approach (bottom left) dilates the sun and reflection and has a strong blurring effect. Neumann et al. (bottom middle) masks details of the background structures, and alters the water’s brightness, giving the impression of another light source. Our approach preserves the lightness of regions, the brightness of the sun, keeps all paint strokes visible, and when visually compared to original contains fewer spatial modifications (bottom right).

Finally, we show results on highly complex images that are more perceptually accurate than similarly sharpened Gimp greyscale. In Figure 10, the hats are more bright and the furthest two are distinguished more easily. In Figure 11, the red fish and stone advance and the two orange fish reappear.




Figure 8: A frame from our hummingbird video. Source: www.naturelibrary.com.

We implemented our conversion in Octave using the Colour Engineering Toolbox and a Matlab toolbox for Laplacian pyramids. Our runtime depends on image resolution and the speed of colour mapping and pyramid construction. We specify times for both 1 and 4 storey pyramids, computed on an Intel4 3 GHz CPU. The Impression Sunrise image (311×223 pixels) takes 1.8 or 3.2 seconds; the Impatiens image (570×593) takes 6.7 or 10.8 seconds; and the Hummingbird video (192×144) single scale conversion takes 136.3 seconds with 0.96 seconds for each of the 142 frames. We also provide a single-scale GIMP plugin at [WWW08].

7. Conclusions and Future Work

In this paper, we have presented a new approach to color to grey conversion. Our approach offers a more perceptually accurate appearance than standard luminance mapping and generates a closer response to the original than other approaches. In particular, we have incorporated the H-K effect which we consider fundamental to obtaining faithful greyscale reproductions. Our two-step approach is a good compromise between a fully automatic technique (first step) and user control (second step) making this approach well suited for natural images, photographs, artistic reproductions as well as business graphics. Finally a major benefit is the consistency we ensure by avoiding changes in color ordering. This makes this technique well adapted to the treatment of videos and image sets.

The main limitation of our approach is the locality of the second step. It can not restore chromatic contrast between non-adjacent regions. This step also risks introducing temporal inconsistencies, which is prevented by constant local parameters. Our algorithm may enhance artifacts from unreliable chromatic information, which occurs in low-quality aggressively compressed media. With respect to colour appearance, we have not taken into account changes of chromatic contrast due to the context of the pixel, specifically its appearance with respect to its surrounding colours.

In future work, we plan to predict image appearance with respect to its color and spatial surrounding by incorporating the masking effect of colour patterns and measuring the *visibility* of the original contrasts using the contrast sensitivity function (CSF) for chromatic channels [Mul85] as a function of spatial frequency. We also plan to investigate methods for modifying enhancement parameters over time while maintaining temporal coherence. Further investigation into the problem of converting video to greyscale or to a reduced colour set is important for video stylization, processing and display on limited devices. We hope our work has fostered new ideas in this direction.

Acknowledgements

We thank Rafal Mantiuk, Hendrik Lensch, Andrei Lintu and the anonymous reviewers for their valuable comments.

Appendix A: Nayatani Chromatic Lightness Model

In CIELUV colour space u', v' are CIE 1976 chromaticity of test stimulus and u'_c, v'_c are chromaticities of the reference white and L_a adapting luminance, set by default to 20 as suggested by Nayatani. Following are equations used in chromatic lightness $L_{N_{VAC}}^*$ and $L_{N_{VCC}}^*$ [Nay97, Nay98].

$$K_{Br} = 0.2717 \frac{6.469 + 6.362L_a^{0.4495}}{6.469 + L_a^{0.4495}} \quad (7)$$

$$s_{uv} = 13[(u' - u'_c)^2 + (v' - v'_c)^2]^{\frac{1}{2}} \quad (8)$$

$$\theta = \tan^{-1} \frac{v' - v'_c}{u' - u'_c} \quad (9)$$

$$q(\theta) = -0.01585 - 0.03017\cos\theta - 0.04556\cos2\theta \\ - 0.02667\cos3\theta - 0.00295\cos4\theta \\ + 0.14592\sin\theta + 0.05084\sin2\theta \\ - 0.01900\sin3\theta - 0.00764\sin4\theta \quad (10)$$

References

- [BA83] BURT P. J., ADELSON E. H.: The Laplacian Pyramid as a compact image code. *IEEE Trans. on Comm.* (1983), 532–540.
- [BB04] BALA R., BRAUN K.: Color-to-grayscale conversion to maintain discriminability. *Proc. SPIE 5293* (2004), 196–202.
- [BE04] BALA R., ESCHBACH R.: Spatial color-to-grayscale transformation preserving chrominance edge information. *Proc. IS&T/SID's 12th Color Imaging Conference* (2004), 82–86.
- [CF03] CALABRIA A., FAIRCHILD M.: Perceived image contrast and observer preference i & ii: The effects of lightness, chroma, and sharpness manipulations on contrast perception. *Journal of Imaging Science & Technology*, 47 (2003), 479–493.
- [Fai05] FAIRCHILD M.: *Color Appearance Models, 2nd Ed.* Wiley-IS and T, Chichester, UK, 2005.
- [FP91] FAIRCHILD M., PIRROTTA E.: Predicting the lightness of chromatic object colors using CIELAB. *Color Research and Application* 16, 6 (1991), 385–393.
- [GD07] GRUNDLAND M., DODGSON N. A.: Decolorize: Fast, contrast enhancing, color to grayscale conversion. *Pattern Recogn.* 40, 11 (2007), 2891–2896.
- [GOTG05] GOOCH A. A., OLSEN S. C., TUMBLIN J., GOOCH B.: Color2gray: salience-preserving color removal. *ACM Trans. Graph.* 24, 3 (2005), 634–639.
- [KMS07] KRAWCZYK G., MYSZKOWSKI K., SEIDEL H.-P.: Contrast restoration by adaptive countershading. In *Computer Graphics Forum (Proc. Eurographics 2007)* (2007), vol. 26.
- [LCD06] LUFT T., COLDITZ C., DEUSSEN O.: Image enhancement by unsharp masking the depth buffer. *ACM Trans. Graph.* 25, 3 (2006), 1206–1213.
- [MMS06] MANTIUK R., MYSZKOWSKI K., SEIDEL H.-P.: A perceptual framework for contrast processing of high dynamic range images. *ACM Transactions on Applied Perception* 3, 3 (July 2006), 286–308.
- [Mul85] MULLEN K. T.: The contrast sensitivity of human color vision to red-green and blue-yellow chromatic gratings. *Journal of Psychology* 359 (1985), 381–400.
- [Nay97] NAYATANI Y.: Simple estimation methods for the Helmholtz-Kohlrausch effect. *Color Res. Appl.* 22 (1997), 385–401.
- [Nay98] NAYATANI Y.: Relations between the two kinds of representation methods in the Helmholtz-Kohlrausch effect. *Color Res. Appl.* 23 (1998), 288–301.
- [NB98] NOWAK R. D., BARANIUK R. G.: Adaptive weighted highpass filters using multiscale analysis. *IEEE Transactions on Image Processing* 7, 7 (1998), 1068 – 1074.
- [NCN07] NEUMANN L., CADIK M., NEMCSICS A.: An efficient perception-based adaptive color to gray transformation. In *Proceedings of Computational Aesthetics 2007* (Banff, Canada, 2007), Eurographics Association, pp. 73–80.
- [RGW05] RASCHE K., GEIST R., WESTALL J.: Re-coloring images for gamuts of lower dimension. *Computer Graphics Forum* 24 (September 2005), 423–432(10).
- [WWW08] WWW: Apparent Greyscale Project <http://www.mpi-inf.mpg.de/resources/ApparentGreyscale/>, 2008.
- [Wys67] WYSZECKI G.: Correlate for lightness in terms of CIE chromaticity coordinates and luminous reflectance. *Journal of the Optical Society of America* (1917-1983) 57 (Feb. 1967).




Figure 9: Our impression is more like the original because it preserves the paint strokes especially in the sky and background. Gooch's image is strongly blurred with a dilated sun and Neumann masks the background and lightens the water.



Figure 10: The extreme brightness of the hats is more apparent in our image than Gimp's greyscale which highlights the differences between the furthest two hats. Source: www.vischeck.com.




Figure 11: Our approach accentuates the red fish and stone, and restores salience to the orange fish (bottom left and right side). Source: Getty Images.