

Contrast Restoration by Adaptive Countershading

Grzegorz Krawczyk, Karol Myszkowski, and Hans-Peter Seidel

Max-Planck-Institut für Informatik
Saarbrücken, Germany

Abstract

We address the problem of communicating contrasts in images degraded with respect to their original due to processing with computer graphics algorithms. Such degradation can happen during the tone mapping of high dynamic range images, or while rendering scenes with low contrast shaders or with poor lighting. Inspired by a family of known perceptual illusions: Craik-O'Brien-Cornsweet, we enhance contrasts by modulating brightness at the edges to create countershading profiles. We generalize unsharp masking by coupling it with a multi-resolution local contrast metric to automatically create the countershading profiles from the sub-band components which are individually adjusted to each corrected feature to best enhance contrast with respect to the reference. Additionally, we employ a visual detection model to assure that our enhancements are not perceived as objectionable halo artifacts. The overall appearance of images remains mostly unchanged and the enhancement is achieved within the available dynamic range. We use our method to post-correct tone mapped images and improve images using their depth information.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation I.4.3 [Image Processing and Computer Vision]: Enhancement

1. Introduction

Successful comprehension of observed images and scenes depends on our ability to distinguish their features. Human vision identifies scene features through the apparent contrasts that they create within their context. Well visible contrasts facilitate the recognition of objects in a scene, identification of their texture, understanding of their spatial distribution, and the ability to judge brightness between adjacent and distant areas. Together, these features directly influence people's assessment of overall image quality [Jan01]. Clearly, a well pronounced rendition of perceived contrasts should be the goal of computer graphics algorithms which process visual information. Unfortunately, often this goal is not achieved due to either technical limitations or poor input data. In tone mapping for instance [RWPD05], the insufficient capabilities of displays require reduction of the dynamic range in images, which inevitably leads to attenuation of contrasts and loss of visual information. In rendering on the other hand, poor design of illumination or bad shading algorithms produce low contrast images in which comprehension of scene content is strongly confined [LCD06].

In this work we are concerned with the problem of communicating contrasts in images that suffered from contrast degradation with respect to their original. In case of a tone mapped image, the original is its source High Dynamic Range (HDR) version. Such HDR images can be captured with HDR cameras, using multi-exposure techniques [RWPD05], or obtained in many rendering applications in particular in realistic image synthesis and lighting simulation. Even if rendering leads to low dynamic range images, e.g. non-photorealistic rendering, contrasts from the depth map can be used for similar purposes [LCD06]. Unlike in typical contrast enhancement tools such as histogram equalization or contrast equalization, we do not want to change the general appearance of processed images. Neither we try to restore the physical contrasts in the image, especially that most often it is not possible due to the dynamic range restrictions. Instead, we propose to enhance the perceived contrasts through a gradual modulation of brightness in the vicinity of the contrasting edge inspired by a family of known perceptual illusions [KM88]: Craik, O'Brien, Cornsweet. These illusions address several models of grad-

ual darkening and lightening of areas towards their common edge to which we in general refer as *countershading profiles*. Our approach has particular advantages that the contrast enhancement can be achieved within the available dynamic range, and the modifications do not change the general appearance of an image because they are limited to the areas along the edges of the enhanced features. Furthermore, the perceived contrast may be larger than would be normally achievable on a target display.

We present an image processing tool that creates countershading profiles for an image to enhance perceived contrast of features degraded with respect to the original. Our tool can be considered as a generalization of *unsharp masking* – an image enhancement technique which in certain cases also creates countershading profiles by overlaying the difference of an image and its blurred version. The development of a new algorithm is motivated by the disadvantages of the traditional unsharp masking which cannot be applied to automatically correct individual image features. To deliver the automatic correction with respect to a reference image, we combine the countershading algorithm with a multi-resolution contrast metric. The metric measures local contrast of features at different scales, compares the processed image to its reference, and drives the spatial extent and the strength of countershading profiles. We first demonstrate how to match the physical amplitude of a reference contrast with the amplitude at the profiled edge, and later we adjust the amplitude according to findings in psychophysics to reduce the perceptual difference between them. Finally, excessive countershading profiles may become visible as halo artifacts and degrade the image quality, what in most cases is unacceptable and in fact reduces the strength of the contrast enhancement. We employ the visual detection model to estimate the maximum amplitude of a countershading profile that is not objectionable in a given area based on the luminance threshold, contrast sensitivity and the contrast masking effects.

We start with a review of unsharp masking and contrast enhancement techniques recently used in computer graphics in Section 2, and we summarize relevant findings in psychophysics in Section 3. Next, in Section 4 we present a new algorithm to create the countershading profiles. In Section 5 we introduce the visual detection model used to adjust the adaptive countershading to prevent undesired halo artifacts, and draft the implementation in Section 6. Finally, we illustrate and discuss possible applications in Section 7.

2. Previous Work

Unsharp masking [Pra91] is the technique in which a Gaussian blurred image Y_σ is subtracted from its original luminance Y to create an unsharp mask that is added to the original image with a coefficient c :

$$\mathcal{Y} = Y + c \cdot (Y - Y_\sigma), \quad (1)$$

where \mathcal{Y} is the enhanced image and σ determines the spatial extent of the Gaussian kernel. The magnitude of the cor-

rection c needs to be adjusted by the user and all pixels in the image are corrected with the same coefficient. However, the enhancement happens in two dissimilar ways: through the countershading and through the reintroduction of features. The highest quality of correction is gained only for image features whose scale is similar to or larger than the size of the Gaussian kernel [Ney93], because they obtain valid countershading profiles. Small kernels, however, lead to sharpening effects at the edges of larger features [Ney93] and have limited capabilities to enhance contrast [KM88]. All features smaller than the kernel size are reintroduced with a varied strength which is influenced by their scale and the difference from the local average, as illustrated in Figure 1. The noise amplification caused by such a reintroduction of the small scale features and the sharpening artifacts at the high contrast edges can be minimized with the adaptive unsharp masking [PRM00, RSMY96]. Psychophysical findings, which show that the uniform physical correction is perceived as stronger in the dark parts of an image than in light areas, motivated the non-linear adaptive unsharp masking [RSMY96]. In spite of the numerous improvements to this technique, we are not aware of any method for an automatic enhancement using individually adjusted kernel sizes and profile magnitudes to create the countershading profiles that are appropriate for enhanced features without distorting other parts of the restored image.




Figure 1: Countershading using unsharp masking gives correct results when the kernel size is adjusted to the size of the feature (left). If unsharp masking is used to enhance the contrast lost on a step edge with details, the filter models the countershading profile on the edge but also strongly amplifies all the features of a smaller scale (right).

The influence of weak contrasts on a limited comprehension of the spatial distribution of objects in a scene have been studied by Luft et al. [LCD06]. They show that unsharp masking using the depth map of a scene strongly enhances the cognition of spatial distribution of objects. Their results are very good because depth maps extract precisely the edges which outline objects in a scene and whose correction improves the perception of the spatial organization. The intensity of countershading, however, depends only on the depth relations of objects behind, and therefore unnaturally looking dark outlines may appear over the objects further behind in the scene. The visual model presented here limits the countershading strength based on the actual image contents to prevent the visible degradations of images, thus limiting such artifacts.




Figure 2: Countershading profile at the borders of areas of equal luminance creates a perceived brightness difference between them.

The loss of communicated information is also typical for tone mapping [RWPD05], where the contrasts are explicitly reduced in an HDR image to fit into the dynamic range of a display or print. Smith et al. [SKMS06] show that, despite the different approaches to tone mapping, each algorithm suffers from a certain amount of contrast degradation leaving space for improvements towards the reference HDR. In their work, to better communicate lost contrast information, fine details are corrected with opposite colors guided by a single-resolution local contrast metric. Then, the largest contrast is restored with a segmentation based countershading technique adjusted by a single global contrast measure. Unfortunately, all features of the intermediate size remain uncorrected and the countershading is applied to only one arbitrary edge in the image. We propose to strongly couple the countershading with a multi-resolution local contrast metric and automatically correct features at various scales in a consistent manner with the individually adjusted profiles. Further, we provide a perception model which counteracts the objectionable halo artifacts.

A comprehensive model of the human visual system is embedded in the Visual Differences Predictor [Dal93], which detects the differences between the reference and distorted images. Such a visual model accounts for luminance masking, spatial contrast sensitivity, and contrast masking in spatial frequency and orientation bands. However, it is computationally expensive and therefore is often simplified in computer graphics applications. Predicting the visible rendering artifacts [RPG99], for instance, is successfully done with a simpler model which ignores the orientation bands. We derive a similar detection model to prevent the countershading profiles from appearing as the halo artifacts. While these models are more focused on the near-threshold noise detection, in our context the suprathreshold effects of lower frequencies are of more interest.

3. Perceptual Background of Countershading

A carefully shaped luminance profile at an edge between two areas, like in Figure 2, causes change in the brightness of the whole areas and increases the perceived contrast between them. Kingdom et al. [KM88] summarize a family of such border profiles and their influence on the brightness of adjacent areas. As shown in Figure 3, practically any form of countershading and the combination of them leads to a

magnification of the perceived contrast. Such profiles can be modelled using the Gaussian function in which the amplitude and the standard deviation determine the intensity of the illusion. Their frequency based modelling algorithm is technically similar to unsharp masking, and leads to alike profiles (compare with Figure 1).




Figure 3: Different countershading profiles and their influence on the perceived brightness. Plots after [KM88].

Dooley and Greenfield [DG77] determined the relation that gives the amplitude of a countershading profile that is required to obtain a perceptual contrast match with a step edge with respect to a simple stimulus. Additionally, Burr [Bur87] observed the increase of perceived contrast when a countershading profile is added to an existing step edge. It has also been found that the spatial extent of a profile determines the maximum possible enhancement and has to be appropriate for the magnitude of the corrected contrast. For instance, a Cornsweat profile of 2 visual degrees can simulate an edge of up to 0.5 Michelson contrast (i.e. strong suprathreshold contrast), but further amplification of this profile leads in fact to a decreased illusion [DG77, Fig.4]. As soon as the low frequency of the profile can be independently detected, the profile is clearly distinguished at an edge and the increase in the amplitude has no effect on the perceived contrast [Bur87]. This suggests that the contrast enhancement using the countershading profiles should be guided by a visual detection model. Finally, the illusions created by the spatially larger profiles are not affected by an additive noise [Bur87], thus the countershading profiles applied to the differently textured areas give consistent effects.

The strength of the perceived contrast enhancement due to countershading is influenced by visual cues. In particular, a contextual hint that the countershading profile results from a difference in the illumination of two surfaces, possibly confirmed by the perspective information, enhances the strength of the effect almost twice [PSL99]. This is confirmed by the success of Luft et al. [LCD06] approach, in which objects separated by different depths are likely to be differently illuminated as well. In contrary, the confidence that a profile is a feature of the surface reflectance significantly reduces the illusion. These observations, especially related to the larger

scale contrasts, cannot be explained by the receptive field properties of the lower order visual neurons, or by the fact that both the step edge and the countershading profile have almost the same frequency characteristics when normalized by the contrast sensitivity function [KM88].

While the early explanations of Cornsweet effect are based on the threshold sensitivity, the illusion is clearly suprathreshold and in fact a consistent theory explaining all experimental findings has not been found so far. Our decision to use the modified suprathreshold sensitivity [DG77] is motivated by the fact that this model explains well the results of the experiments which measure the apparent contrast, including the suprathreshold effects, of up to 0.7 in Michelson measure. Clearly, the visual cues strongly articulate the effect [PSL99], but even if an appropriate model was available, it would require a robust decomposition into illumination and reflectance which practically is only possible in rendering.

4. Image Processing for Countershading

We develop a method that creates the countershading profiles to enhance the perceived contrasts of edges in the restored (input) image that are less pronounced than the corresponding contrasts in the specified reference image. We identify such edges in the restored image by comparing it to the reference image using the multi-resolution local contrast metric. Guided by the metric, we create the profiles from the sub-band components such that the profiles are individually adjusted to the corrected features without distorting information that has been well preserved from the reference.

4.1. Multi-resolution Local Contrast Metric

We use a metric which measures the physical local contrast at several frequency bands in a similar manner to Peli [Pel90]. We decompose an image into the Gaussian pyramid, in which each lower level is filtered by a 5×5 kernel and its resolution is reduced twice as described in [BA83]. Such a decomposition splits the image frequencies into octaves what corresponds to the frequency separation observed for the human visual system [Pel90]. Thus, on the highest level we measure the contrast of fine details, and on the lowest level the contrasts between the major areas in the image. The lowest level we consider is 4 pixels long in the smaller dimension, and we ignore the base band. For each pixel at each pyramid level l , we calculate the local sub-band contrast C_l using the formula:

$$C_l = \frac{|Y - Y_{mean}|}{Y_{mean}}, \quad (2)$$

where Y is the luminance of a pixel at the pyramid level l and Y_{mean} is the local mean luminance. In practice Y_{mean} is taken from the corresponding pixel at the lower pyramid level. The final output of the metric is the pyramid that contains the ratios of the corresponding local contrasts between

the restored (input) image and its reference:

$$R_l = \frac{C_l^{inp}}{C_l^{ref}}. \quad (3)$$

The ratio R_l is limited to the maximum value of 1 because the detail amplification with respect to the reference is not considered.

4.2. Adaptive Countershading

We develop a method which selectively adds the countershading profiles to the restored image guided by the sub-band local contrast ratio (3) from the metric. We start with an observation that the addition of successive levels of the full resolution Difference of Gaussians up to a certain level (here the example for 3 levels):

$$U = (Y - Y_{\sigma(1)}) + (Y_{\sigma(1)} - Y_{\sigma(2)}) + (Y_{\sigma(2)} - Y_{\sigma(3)}), \quad (4)$$

gives the same result as unsharp mask, equation (1), for this level: $U = (Y - Y_{\sigma(3)})$, where $\sigma(l) = 2^{l-1}/\sqrt{2}$ denotes the Gaussian blur at the level l and Y is the luminance of the reference image. When such a sum is further informed by the multi-resolution metric which locally adjusts the amplitudes of the sub-band components:

$$P = \sum_{l=1}^N (1 - \uparrow R_l) \times (\log Y_{\sigma(l-1)}^{ref} - \log Y_{\sigma(l)}^{ref}), \quad (5)$$

we obtain the countershading profiles P which are adjusted to match the contrasts in the reference image. In equation (5) l denotes the level of the Gaussian pyramid with N being the lowest, R_l are the contrast ratios from the metric at the selected level, operator (\uparrow) denotes upsampling to the full resolution, operator \times is the element-wise multiplication, and $Y_{\sigma(0)}^{ref}$ is the luminance of the reference image. The difference $(\log Y_{\sigma(l-1)}^{ref} - \log Y_{\sigma(l)}^{ref})$ is a sub-band component of the countershading profile at the level l . The luminance Y and the countershading profiles P are calculated in the logarithmic space. Such a coarse approximation of brightness prevents too strong darkening which would happen in the linear space. The sub-band components are not taken from the contrast metric, but are stored in the full resolution in order to preserve the phase information which would be lost by the resolution reduction. The contrasts in the input image are restored by adding the countershading profiles P to the luminance of the input image in the logarithmic space.

Equation (5) leads to good results for several reasons. The uncontrolled amplification of features do not happen because the algorithm is adjusted to the reference image. The countershading profiles are created from the reference image, because certain features might have been lost in the input image, thus both detail enhancement and detail reintroduction are solved using one framework. The sharp edges of large scale features are detected by the contrast metric on the top level and on all the levels down to the scale corresponding to the size of these features, therefore the countershading for

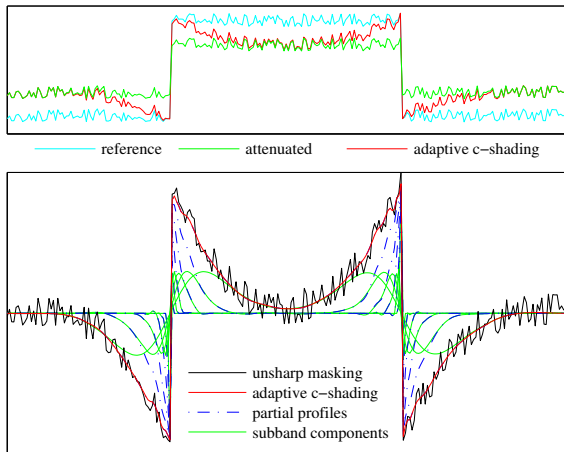


Figure 4: Countershading profile for a step edge with details (top), where the step edge is attenuated while the details are preserved with respect to the reference. Adaptive countershading (bottom) recovers the smooth profile which prevents artifacts. The unsharp mask profile is distorted by the high frequency contents of the reference edge and exaggerates details during enhancement as shown in Figure 1. Unlike in our method, unsharp masking also requires manual adjustment of the spatial extent and the amplitude of the profile.

such edges is progressively composed from the sharp components and the components with a larger spatial extent. This is illustrated in Figure 4 along with the comparison to the traditional unsharp masking.

At this stage, the multi-resolution contrast metric assures that the physical contrast of the features in the image restored with the countershading profiles are equal to their physical contrast in the reference image.

4.3. Saturation of Profiles

Countershading profiles may increase or decrease luminance values beyond the available dynamic range and cause the saturation to black or white, removal of details, and clearly reveal the presence of a profile. Therefore, the parts of the profile that correct beyond the available range have to be attenuated, as shown in Figure 5. The attenuation is performed successively starting from the lowest frequency sub-bands, and separately for the darkening and the lightening parts of the profiles. Each sub-band component is attenuated so that the restored image plus the countershading profiles does not exceed the dynamic range. We motivate our bottom-up approach with the fact that the saturation is mostly caused by the much larger amplitudes of the low-frequency components. Obviously, the strength of the contrast enhancement is reduced in such case, still the asymmetric profiles increase the perceived contrast, as shown in Figure 3, and the degradation of the restored image is prevented.




Figure 5: Countershading may exceed the dynamic range limit and cause clipping: the profile is then distorted and fine details are removed (left). Successive attenuation of a profile retains all details and as much of the profile as possible (right).

5. Perception of Countershading Profiles

The countershading profiles modulate physical contrasts at edges in an image in order to increase the perceived contrasts between features. However, as soon as the low frequency of a profile can be independently detected, the whole profile is distinguishable at an edge, and the increase in the profile amplitude has no further effect on the perceived contrast [Bur87]. To counteract such situations, we develop a visual detection model which assures that the sub-band components of profiles remain below the objectionable amplitude.

We use a model which explains the behavior of the Cornsweet illusion with a good accordance to the perceptual experiments which match the apparent contrast of a profile with the contrast of a step edge [DG77]. The model is based on a spatial contrast sensitivity function (CSF), but its sensitivity to the low frequencies varies with the amplitude of a profile, as shown in Figure 6. It therefore estimates the amplitude thresholds above which the components of the profile become individually visible and render a much weaker Cornsweet illusion with objectionable halo artifacts. This model, however, analyzes single Cornsweet profiles on a uniform 2D background and it may be too conservative for natural images. The contrast masking effect [Wan95] suggests that in the areas which already contain features of certain spatial and orientation characteristics, the acceptable amplitude of the profile may be higher if the profile is composed of signals with the similar characteristics. Such a selectivity of independent visual channels fits well to our multi-resolution contrast analysis which uses filter banks motivated by the visual channels in human perception. We therefore improve the model by accounting for this strong effect in human perception.

An important insight from [DG77], shown in Figure 6, is that the sensitivity to the higher frequencies in the Cornsweet profile and in the step edge is similar which justifies our approach to equal the contrast at the profile edge to the contrast of the feature edge.

We take a standard approach to modeling visual detection models [Dal93] in which we combine three effects typical to the human vision: luminance masking, spatial contrast sen-

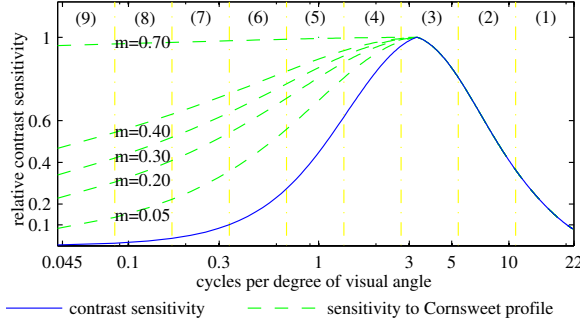


Figure 6: Contrast sensitivity function for the threshold effects and the supra-threshold model of tolerance to the magnitude of sub-band components of the Cornsweet profile added to an existing step edge of Michelson contrast m . Vertical lines denote frequency ranges of sub-bands at pyramid levels numbered in the top.

sitivity and contrast masking. For the sub-band component at level l of our pyramid representation, the maximum amplitude of a profile ΔY expressed in luminance is calculated as follows:

$$\Delta Y = \frac{tvi(Y_{mean})}{csf_l} \cdot \text{masking}(C_l, Y_{mean}), \quad (6)$$

where Y_{mean} is the local mean luminance in the sub-band (considered as the adapting luminance) and C_l is the sub-band contrast at the level l . tvi is the threshold versus intensity function [CIE81], csf_l is the relative loss of contrast sensitivity for the spatial frequency band l [Dal93], and masking describes the contrast masking at the given contrast C_l . The ΔY is calculated for each pixel at the sub-band level l .

The threshold versus intensity function tvi [CIE81] describes the luminance masking effect by giving the minimum luminance change which is visible at the adapting luminance level [Wan95]. While tvi describes the thresholds for a patch shown on a uniform background, the human vision response to complex images varies depending on the frequencies of the components. Our sensitivity to the contrast at a given frequency is described by the contrast sensitivity function (CSF), which in practice increases the luminance thresholds estimated by tvi for very high and low frequencies. In our model, however, instead of using the typical CSF which describes the detection thresholds, we use the function given in [DG77, Fig.6] which determines the tolerable amplitudes of the countershading profiles and also accounts for the increase in sensitivity to low frequencies when a profile is added to an existing edge of Michelson contrast m . Due to the lack of equation, we provide a fit based on the normalized CSF [Dal93]:

$$csf_l(m) = csf_l^{0.74 - 0.83 \cdot m^{0.35}} \quad (7)$$

This function replaces csf_l in equation (6) for levels l with frequencies lower than frequency of peak sensitivity $\approx 5cpd$. Since in the original publication this relation is expressed

using Michelson contrast, for compatibility we recalculate here our contrast measure C . The plot is given in Figure 6.

Signals added to textured areas are harder to perceive than if added to uniform areas. This happens because existing contrasts in an area mask the contrast of the introduced signal. Contrast masking elevates the detection threshold as a function of the local sub-band contrast C_l in the corrected image:

$$\text{masking}(C_l, Y_{mean}) = \max \left(1, \left(\frac{C_l}{T_C(Y_{mean})} \right)^{0.7} \right), \quad (8)$$

where T_C is the threshold contrast for the local mean luminance level Y_{mean} , and $T_C(Y_{mean}) = \frac{tvi(Y_{mean})}{Y_{mean}}$. Contrast masking is modeled by a power function with a typical exponent 0.7 [Dal93], which increases the thresholds as soon as the local sub-band contrast is greater than the threshold contrast.

Contrast masking is normally considered within the frequency band and the orientation band [Dal93]. We ignore the orientation bands due to the high computational costs. In case of the low frequencies, the introduced profile in our case has the same orientation as the existing signal (corrected edge) which gives a strong masking effect.

The maximum tolerable amplitude of the profile ΔY from equation (6) sets the limit for the amplitude of the sub-band component of the countershading profile at the given location.

5.1. Natural Image Statistics

One aspect evident in the analysis of Cornsweet [DG77], is that strong contrasts cannot be corrected with small profiles. However, according to findings in natural image statistics [BGB00], the average amplitudes of frequencies in images tend to decay as a power function, being large for the low frequencies and small for the high frequencies. Such a phenomenon is known as the power law for the amplitude of frequencies. This observation assures that in the context of natural images, we are highly unlikely to encounter the correction of a very high contrast step with a small profile.

6. Implementation

The adaptive countershading algorithm restores the degraded image Y^{imp} with respect to its reference Y^{ref} and outputs an enhanced version of Y^{imp} . The algorithm operates only on luminance values. To process a color image, the RGB channels are converted to Yxy color space and reverted back to RGB with an enhanced luminance channel Y. We clip the colors that are mapped out of the sRGB gamut after processing. Nonstandard references, such as depth maps, may be directly used instead of contrast ratios R . However, such ratios have to be manually scaled to reasonably guide the strength of the contrast enhancement. We give the algorithm outline below:

```

Cinp = contrasts_pyramid(Yinp);
Cref = contrasts_pyramid(Yref);

Pc = profile_components(log(Yref));
P = 0; // countershading profile

n = log2( min(width,height) );
for l:=n..1
  R(l) = Cinp(l)/Cref(l);
  aR = 1-min(1,R(l));
  aS = saturation_limit(log(Yinp)+P,Pc(l));
  aV = visual_model_limit(Yinp,Cinp(l),Pc(l));
  P += Pc(l) * upsample( min(aR,aS,aV) );

RESULT = 10^(log(Yinp) + P);

```

The process is fully automatic given the reference image Y^{ref} or arbitrary data passed as the contrast ratios R . The visual model assumes an sRGB display and requires that the image frequencies are calculated in cycles per degree of visual angle which depend on the screen resolution and the viewing distance. The results in Section 7 are obtained for the resolution 1280×1024 viewed from the distance of $1.5 \times$ the screen height. An enhancement of a 1Mpx image requires about a minute on a modern PC. The bottleneck of the algorithm are the convolutions, three are calculated per pyramid level: to measure the contrasts in Y^{ref} , contrasts in Y^{inp} , and to calculate the components of P . A linear filter is used for upsampling.

7. Results and Applications

We first demonstrate adaptive countershading on a test pattern, Figure 9. The reference image (a) contains a textured background and two textured patches. After the tone mapping (b), the texture of the right patch has been preserved, while the textures of the background and the left patch have been attenuated. Also, the contrast between both patches and background has been attenuated. Thus the left patch illustrates global tone mapping and the right one local. The goal of the correction is to restore the contrast between the patches and the background, to restore the visibility of the textures in the background and the left patch, to assure that the texture of the right patch is not emphasized and that the objectionable halo artifacts do not appear. The image (a) spans the full dynamic range and in the images (b,c,e) the dynamic range is artificially limited for the demonstration purposes. The countershading (c) visibly enhances the contrasts comparing to the tone mapped image (b). The texture details of the background and the left patch are restored to almost the same level as in the reference image (a). The contrast and the brightness of the right patch has also improved, although it cannot match the reference due to the dynamic range restrictions. The details of the right patch remain unchanged, which is confirmed in the map (d). Unsharp masking with the spatial extent and the magnitude manually adjusted for correction of the patch to background contrast is shown in image (e). The image is visibly enhanced, however, when compared to the reference (a), the background and the

right patch details have clearly too strong magnitude. The undesired halo is well visible in the right patch where also some areas became saturated.

7.1. Post Tone Mapping Restoration

In tone mapping [RWPD05], the insufficient capabilities of displays require the reduction of dynamic range in images, which inevitably leads to the attenuation of contrasts and loss of visual information with respect to the original HDR. In Figure 7 an HDR image has been tone mapped with a contrast equalization technique [MMS06] to reveal the details. Unfortunately, the result does not depict any more the strong brightness difference between the clouds and the building which is very apparent in the original image. This has been detected by the multi-resolution contrast metric and corrected with the appropriate countershading profiles to reintroduce the brightness difference. After the enhancement, the overall appearance of the tone mapped image including the fine details is not changed. Such a correction is not possible with unsharp masking, although the size and the magnitude of the blur in the mask has been manually adjusted according to the metric data. The reason is that the larger kernel, which is required for this correction, amplifies details so strong that the countershading effect disappears. Another example is shown in Figure 8, where an HDR image has been tone mapped with the logarithmic mapping [DMAC03]. After using this global operator, some cloud details in the sky are not visible any more, the area around the sun becomes almost isoluminant, and much contrast has been lost in the horizon area. This is automatically restored with the adaptive countershading and the style of the particular tone mapping algorithm is not changed. In both examples the halo artifacts are not disturbing even though a stronger correction was allowed by the visual detection model in Figure 7 because of the masking by the clouds.

7.2. Adaptive Depth Sharpening

Unsharp masking using the depth map of a scene strongly enhances cognition of spatial distribution of objects [LCD06]. We obtain a similar enhancement using the adaptive countershading by measuring the relation of a depth map of an image to a uniform map in place of the contrast ratios R in equation (3) and by using the depth map instead of the reference luminance in equation (5). In our approach, Figure 10, the intensity of countershading does not only depend on the depth relations of objects, but is also guided by the visual model which prevents the appearance of unnaturally looking dark outlines over objects further behind in the scene. The visual model limits the countershading strength based on the actual image contents and prevents visible degradations.

8. Conclusions

Based on findings from psychophysics, the paper explains how to enhance contrast in images using the Craik-O'Brien-Cornsweet illusion in a controlled way by employing the



Figure 7: Image tone mapped using the contrast equalization [MMS06] (top/left) and restored by adaptive countershading (bottom/left). The restored image better communicates the brightness relations and the depth in the image. (top/right) shows unsharp masking with parameters set manually to equal dominant countershading profile. Although overall enhancement of unsharp masking is impressive, the changes are hardly controllable and modify the style of the image.

multi-resolution local contrast metric to guide the strength of enhancement and the visual detection model to prevent the appearance of objectionable artifacts. Countershading in most cases cannot be expected to restore the original contrast of the reference, however the enhancement is well visible when profiles are well adjusted and are masked by image contents.

We present an image processing tool to create countershading profiles which are individually and automatically adjusted to enhance selected image features that require such correction when compared to the reference. The same framework is also able to reintroduce lost contrast information. We demonstrate how it can be used to enhance images using their HDR originals or the depth information as the reference. Comparing to the results of the traditional unsharp masking, the enhanced images better communicate information through contrast while the overall appearance is not distorted and the enhancement is achieved within the available dynamic range.

We would like to pursue this research and evaluate the achieved corrections with perceptual experiments, which measure the actually perceived strength of the countershad-

ing enhancement in complex images for stimuli of different scales and given a variety of contrast references.

Acknowledgements

We would like to thank Rafał Mantiuk and the anonymous reviewers for their valuable comments concerning this work.

References

- [BA83] BURT P. J., ADELSON E. H.: The Laplacian pyramid as a compact image code. *IEEE Transactions on Communications COM-31*, 4 (April 1983), 532–540. 4
- [BGB00] BOVIK A. C., GIBSON J. D., BOVIK A. (Eds.): *Handbook of Image and Video Processing*. Academic Press, Inc., Orlando, FL, USA, 2000. 6
- [Bur87] BURR D.: Implications of the Craik-O’Brien illusion for brightness perception. *Vision Research* 27, 11 (1987), 1903–1913. 3, 5
- [CIE81] CIE: *An Analytical Model for Describing the Influence of Lighting Parameters Upon Visual Performance*, vol. 1. Technical Foundations, CIE 19/2.1. International Organization for Standardization, 1981. 6



Figure 8: Image tone mapped using logarithmic mapping [DMAC03] (top) and restored using adaptive countershading (middle). Subtle changes to the image bring back the contrast at the horizon and the details, but do not change the style of the image. Image courtesy of SpheronVR.

- [Dal93] DALY S.: The visible differences predictor: An algorithm for the assessment of image fidelity. In *Digital Images and Human Vision* (1993), Watson A. B., (Ed.), MIT Press, pp. 179–206. ISBN: 0-262-23171-9. 3, 5, 6
- [DG77] DOOLEY R. P., GREENFIELD M. I.: Measurements of edge-induced visual contrast and a spatial-frequency interaction of the Cornsweet illusion. *Journal of the Optical Society of America* 67 (1977). 3, 4, 5, 6
- [DMAC03] DRAGO F., MYSZKOWSKI K., ANNET T., CHIBA N.: Adaptive logarithmic mapping for displaying high contrast scenes. In *Proc. of Eurographics* (2003), Brunet P., Fellner D., (Eds.), pp. 419–426. 7, 9
- [Jan01] JANSSEN R.: *Computational Image Quality*. SPIE Press, Bellingham, WA 97227-0010, USA, 2001. ISBN 0-8194-4132-5. 1
- [KM88] KINGDOM F., MOULDEN B.: Border effects on brightness: a review of findings, models and issues. *Spatial Vision* 3, 4 (1988), 225 – 262. 1, 2, 3, 4
- [LCD06] LUFT T., COLDITZ C., DEUSSEN O.: Image enhancement by unsharp masking the depth buffer. *ACM Transactions on Graphics* 25 (2006), 1206–1213. 1, 2, 3, 7, 10
- [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 (2006), pp. 286 – 308. 7, 8
- [Ney93] NEYCENSAC F.: Contrast enhancement using the laplacian-of-a-gaussian filter. *CVGIP: Graph. Models Image Process.* 55, 6 (1993), 447–463. 2
- [Pel90] PELI E.: Contrast in complex images. *Journal of the Optical Society of America A* 7 (Oct. 1990), 2032–2040. 4
- [Pra91] PRATT W. K.: *Digital image processing (2nd ed.)*. John Wiley & Sons, Inc., New York, USA, 1991. 2
- [PRM00] POLESEL A., RAMPONI G., MATHEWS V.: Image enhancement via adaptive unsharp masking. *IEEE Transactions on Image Processing* 9 (2000), 505–510. 2
- [PSL99] PURVES D., SHIMPI A., LOTTO B. R.: An empirical explanation of the Cornsweet effect. *J. Neurosci.* 19, 19 (1999), 8542–8551. 3, 4
- [RPG99] RAMASUBRAMANIAN M., PATTANAIK S. N., GREENBERG D. P.: A perceptually based physical error metric for realistic image synthesis. In *Proc. of ACM SIGGRAPH 1999* (1999), ACM Press, pp. 73–82. 3
- [RSMY96] RAMPONI G., STROBEL N., MITRA S. K., YU T.-H.: Nonlinear unsharp masking methods for image contrast enhancement. *Journal of Electronic Imaging* 5 (July 1996), 353–366. 2
- [RWPD05] REINHARD E., WARD G., PATTANAIK S., DEBEVEC P.: *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Morgan Kaufmann, 2005. 1, 2, 7
- [SKMS06] SMITH K., KRAWCZYK G., MYSZKOWSKI K., SEIDEL H.-P.: Beyond tone mapping: Enhanced depiction of tone mapped HDR images. In *Proc. of Eurographics 2006* (2006). 3
- [SS03] SCHARSTEIN D., SZELISKI R.: High-accuracy stereo depth maps using structured light. *CVPR* (2003), 195. 10
- [Wan95] WANDELL B. A.: *Foundations of Vision*. Sinauer Associates, Sunderland, Massachusetts, 1995. 5, 6

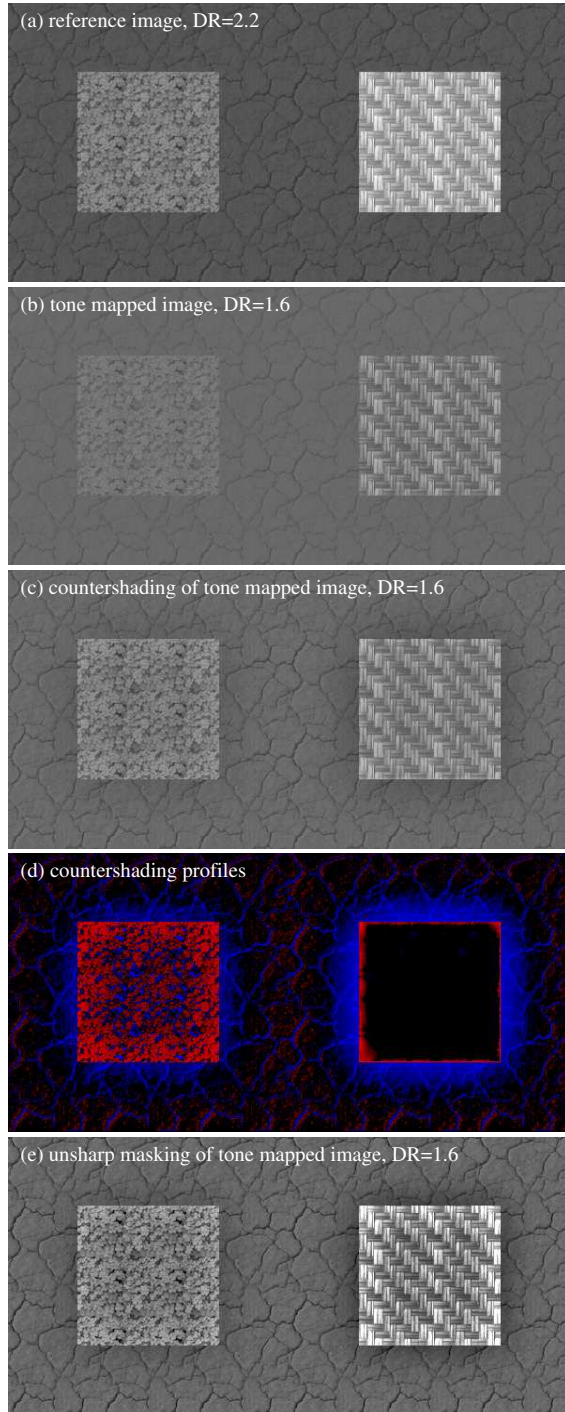


Figure 9: Test pattern for the contrast restoration by adaptive countershading. DR describes the dynamic range in \log_{10} units of luminance. In image (d), the blue countershading profiles darken the image and red lighten, their intensity corresponds to the profile magnitude. The right patch in (d) obtained no lightening because of the dynamic range limit. Refer to Section 7 for details.




Figure 10: Countershading using depth information (a) enhances cognition of the spatial distribution of objects in the scene (b). The visual models limits the appearance of countershading as halo artifacts. Unnaturally looking dark outlines may appear over objects further behind in the scene if only depth relations are considered, image (c) from [LCD06]. Image and depth data from [SS03].