Image Enhancement by Unsharp Masking the Depth Buffer

Thomas Luft* Carsten Colditz* Oliver Deussen*
University of Konstanz, Germany

Figure 1: Drawings by S. Dali, P. Picasso, G. Seurat, and H. Matisse. The artists separate objects of different depth by locally altering the contrast in order to enhance their depth perception. In our work we mimic such an effect for computer generated images and photographs that contain depth information.

Abstract

We present a simple and efficient method to enhance the perceptual quality of images that contain depth information. Similar to an unsharp mask, the difference between the original depth buffer content and a low-pass filtered copy is utilized to determine information about spatially important areas in a scene. Based on this information we locally enhance the contrast, color, and other parameters of the image. Our technique aims at improving the perception of complex scenes by introducing additional depth cues. The idea is motivated by artwork and findings in the field of neurology, and can be applied to images of any kind, ranging from complex landscape data and technical artifacts, to volume rendering, photograph, and video with depth information.

CR Categories: I.3.3 [Computer Graphics]: Picture/Image Generation—Display algorithms; Bitmap and framebuffer operations; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

Keywords: artistic tone reproduction, non-photorealistic rendering, image enhancement, complex scenes

1 Introduction

The capabilities of available modeling and animation software as well as the powerful performance of today's computers, allows also non-computer graphic artists to produce and render scenes containing millions of triangles. However, using the standard shading technique that is provided by graphics hardware, the user cannot obtain the desired quality, since the rendering of complex spatial arrangements sometimes suffers from a dull appearance. This is especially true for shadowed and ambient lighted areas offering only a limited depth perception. We present a method to enhance the perceptual quality of such images by explicitly utilizing the depth information. Although, our idea is motivated by artwork, we do not describe a rendering method solely applicable for non-photorealistic rendering. Rather, since for many images depth information is already available today, our approach can be utilized to enhance all kinds of synthetic images, especially complex scenes in the fields of volume rendering, technical imagery, landscapes, photographs, and videos containing depth information.

Artists tend to enhance contrast and alter color locally to support the perception of objects in an image. Figure 1 shows some artwork that demonstrates the effect. Near the boundary of objects that are located in front of other objects the tone is locally altered: the background or partly occluded objects are slightly darkened to create the impression of depth or bright halos are placed around objects to visually detach them from the background. Such drawing style could be motivated by findings in the field of neurology. For example, we known from cognitive psychology that the retina of the eye locally intensifies the contrast to enhance the recognition of objects (see [Eysenck and Keane 2000]). Consequently, when artists emphasize a contour and alter the local contrast, they additionally support the viewing process. This behavioral mechanism is also relevant in tone mapping. Similar to computer imagery, artists have to represent a scene with potentially high contrast on a medium with a low dynamic range. In contrast to tone mapping algorithms, the painter has a concrete idea of the scene's depth, and is therefore able to visually encode the spatial relationships between the objects.

We present a computational method to achieve similar effects for images that contain depth information. Our method is closely related to image space techniques that produce a local contrast enhancement, e.g. the Laplacian-of-a-Gaussian filter [Neycenssac 1993] or the popular unsharp mask technique [McHugh 2005]. In contrast to those methods, we perform depth buffer operations to obtain additional information about spatially interesting areas. This information is then used to alter the input image, and in turn introduces additional depth cues. Applying our technique, we are able to perform a color and contrast enhancement that is modulated by the spatial relations of the objects in a scene.

*[luft,colditz,deussen]@inf.uni-konstanz.de

The paper is structured as follows: after a discussion of the relevant literature in Section 2, we present our algorithm and describe its relation to other known techniques in Section 3. Following, we introduce a number of possible effects and show their application on complex computer-generated objects, rendered global illumination images, and photographs and videos with depth information. We discuss our approach in Section 4, and conclude the paper with an outline of future work.

2 Related Work

Our work is related to several research fields. As mentioned above, tone mapping of high dynamic range images involves a compression of local contrast, which has to be implemented in a perceptually feasible way. We also find several image manipulation techniques to enhance the local contrast such as the unsharp mask. In the field of non-photorealistic rendering, image enhancement and abstraction is of special interest, and several methods were presented to alter color and tone of the rendered images in order to introduce perceptional cues for the recognition of additional spatial information.

Tone Mapping and Image Enhancement: One of the widely accepted assumptions for tone reproduction states that the human eye is mostly sensitive to local contrast and not so much to absolute luminance (see [DiCarlo and Wandell 2000]).

Reinhard et al. [2002] present a tone mapping procedure on the basis of the Zone system introduced by Adams [1980]. In their work a local contrast enhancement is applied in addition to global operations. For special values, this operator produces unwanted ringing artifacts around objects (see Figure 9 in [Reinhard et al. 2002]), though, in combination with depth information such effects can be useful for spatial enhancement. For an overview of different tone mapping operators and a user study we refer to Ledda et al. [2005] and Reinhard and Devlin [2005].

Approaches regarding image enhancement can be divided into global and local methods. Global methods mainly represent histogram modifications, see e.g. [Hummel 1975; Stark 2000]. These techniques aim to exploit the full dynamic range of a rendering device by modifying the histogram of an image or parts of it. The attractiveness is their simplicity and the minor computational effort. However, in context of our technique the local methods are of particular interest, since they allow the user to modify the contrast within the vicinity of a certain point. Especially local contrast enhancement based on edge detection filters are widely used. An introduction to this topic is given by McHugh [2005], and related works are proposed by Neycenssac [1993] and Beghdadi and Negrate [1989]. A wavelet based approach for local contrast enhancement is introduced by Starck et al. [2003]. Furthermore, color vision models such as the Retinex model that is also applied for tone mapping, were used for locally amplifying contrasts, see e.g. [Meylan and Süsstrunk 2004]. Nevertheless, the missing depth information limits the possible effects. Our technique can be combined with all of these approaches, since we modify the tone of the input image as well, while additionally consider the depth information.

Non-photorealistic Rendering: Saito and Takahashi [1990] proposed the effect of local enhancement of images using depth information. The authors improved the rendering of an object by utilizing the depth buffer and its derivatives for the creation of contours, creases, and hatching strokes. This work was a major inspiration for us, although, derivatives of the depth buffer are not optimal for visually separating objects in a complex scene. Here many discontinuities of the depth function are found that cause peaks in the derivatives. In contrast to their method, we use the information that can be obtained from the difference between the original and the low-pass filtered depth buffer to modulate contrast and colors.

An artistic contrast enhancement was used by Winkenbach et al. [1994]. They partly removed texture details of surfaces for artistic reasons and for a better perception of synthetic drawings. This kind of contrast enhancement might also be a motivation for the various toon and silhouette rendering techniques that were presented in recent years.

A technique that modulates the amount of visual details to represent objects in a complex botanical scene was presented by Deussen and Strothotte [2000]. The differences in the depth buffer were used to guide the silhouette rendering. If the depth differences exceed a given threshold, a line is drawn. Thus, the spatial importance is depicted: if objects have a larger depth distance, it is more important to draw the separating border. Our work is closely related to their method, since we also consider depth differences to emphasize complex spatial arrangements.

Another important work is the *Non-photorealistic Camera* introduced by Raskar et al. [2004]. Their system produces illustrations, precisely silhouette drawings, of photographs using a multi-flash camera, and thus, also allows to enhance a photograph with additional depth cues.

Gooch et al. [1998] present a non-photorealistic shading model that aims to enhance the image quality in regions that appear dull using standard rendering techniques. In combination with silhouette lines, a variety of drawing styles can be produced that especially improve the perception of dark and ambient areas. While their technique allows altering the appearance of single objects, our technique works on sets of objects and their spatial separation. Nevertheless, we can apply their color model to change colors along depth discontinuities. A comparable work related to shading algorithms is introduced by Cignoni et al. [2005]. This work modifies the surface normals in order to alter the shading of objects. This way, edges are emphasized and surfaces are slightly shaded to likewise diminish especially the dull appearance produced by standard shading procedures.

Another method that was recently used for enhancing the display of spatially complex scenes by providing additional contrast is the computation of ambient occlusions [Pharr and Green 2004; Bunnell 2005]. The ambient occlusion determines the percentage of the hemisphere above an object point that is not occluded by other parts of the object. For each part of the object an additional texture is determined that especially darkens concave surface areas. While theirs is an object space method, our method works in screen space and therefore adapts to changes of the camera, e.g. the changes are getting smaller within a camera zoom. This allows to enhance even small details in a zoomed view. Furthermore, we do not need to precompute the textures and are able to introduce a variety of effects into our scenes. In Section 3.2 the methods are compared.

3 Unsharp Masking the Depth Buffer

This section illustrates the relevance of integrating depth information into the shading process. We firstly formulate several shading styles as functions of the depth, as shown by the example scene in Figure 2. This simple scene contains two objects, a sphere and a cube, which are arranged in front of each other. The depth buffer values and the resulting luminance values along a horizontal scanline are shown in Figure 2(a) and Figure 2(b). A silhouette drawing of the scene looks similar to Figure 2(c), and resembles a bi-

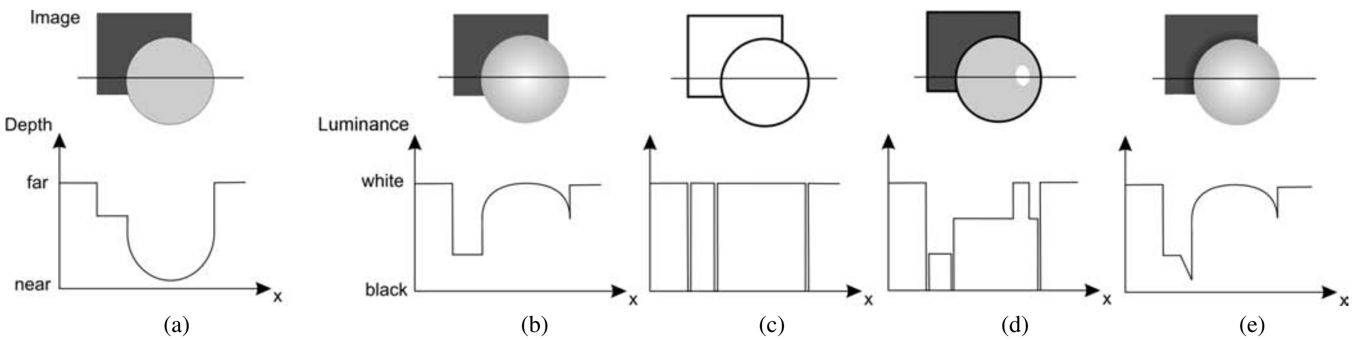

Figure 2: Example scene: (a) depth buffer along a scanline; (b-e) shading functions along a scanline: (b) Phong shading); (c) silhouette rendering; (d) toon shading; (e) haloed contour.

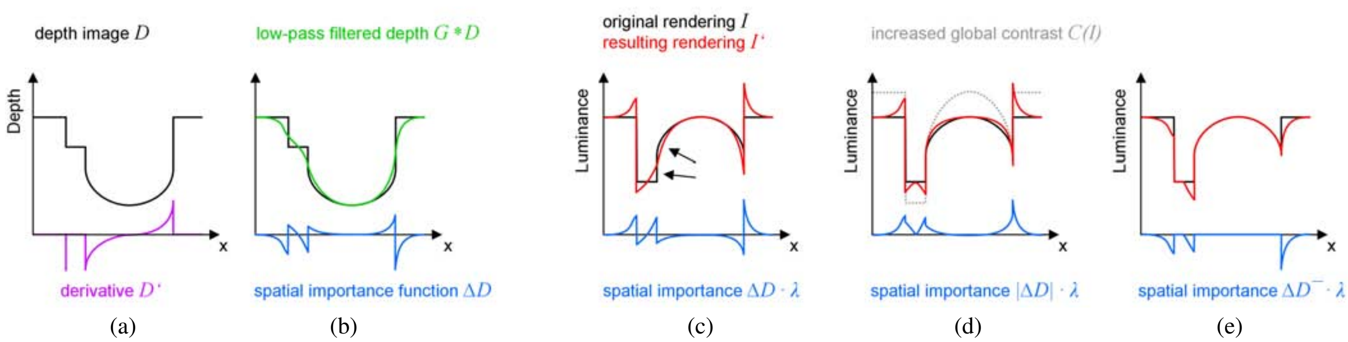

Figure 3: Depth functions: (a) depth function and its derivative; (b) difference between original and low-pass filtered depth buffer. Resulting shading: (c) adding the spatial importance function $\Delta D \cdot \lambda$ with $\lambda < 0$; (d) linear combination of the original input image and a high contrast version weighted by the absolute value of the spatial importance function $| \Delta D | \cdot \lambda$ with $\lambda > 0$; (e) depth darkening by using the negative fraction of the spatial importance function $\Delta D^{-} \cdot \lambda$ with $\lambda > 0$, which proved to be a natural image enhancement operator.

nary representation of the depth buffer derivative. Toon shading as shown in Figure 2(d) is a combination of the silhouette rendering and a quantized color representation. To achieve effects with a larger range such as the haloed lines in Figure 2(e) or the effects motivated by the artworks of Figure 1, we have to obtain information about the spatial relations between the objects.

Accordingly, we propose a technique that is similar to the local contrast enhancement using an unsharp mask. Here, a mask is created by subtracting a low-pass filtered copy from the input image, which effectively results in a high-pass filter. The resulting high frequency signals are added to the original image to obtain a sharpening or a local contrast enhancement. Especially for the latter effect, a larger filter kernel size is used resulting in an enlarged radius of influence. The result is a smooth image modification distributed over a larger area. This technique works with conventional image data without considering depth information, and as a result, enhances all local contrasts in an image. However, by applying this technique to the depth buffer—i.e. we compute the difference between the original and the low-pass filtered depth buffer—we convey an additional semantic information: the spatial relation between the objects in the scene. For example, by applying a conventional unsharp mask to an image that contains a blue object in front of a blue background, the image cannot be altered due to the lack of color contrast. Our technique enables us to still alter the colors in this scenario due to the additional depth information.

Figure 3 demonstrates our approach using the scene of Figure 2. The depth buffer and its derivative is shown in Figure 3(a). While the derivative of the depth function can be used for indicating silhouette lines, it fails to generate halos and other effects due to its very local nature. Consequently, it is not appropriate for the visual depth differentiation of complex scenes with many discontinuities of the depth buffer (see also [Deussen and Strothotte 2000]). This is the reason for proposing a technique similar to the unsharp mask. Therefore, we compute the difference between the original and the low-pass filtered depth buffer as shown in Figure 3(b). We call the resulting function the *spatial importance function* ΔD. Given the input image $I : \mathbb{N}^2 \rightarrow \mathbb{R}^3$ within the RGB color space and the corresponding depth buffer $D : \mathbb{N}^2 \rightarrow \mathbb{R}$ with $d(x,y) \in [0,1]$, we compute ΔD by

$$\Delta D = G * D - D \tag{1}$$

with $G * D$ being the convolution of a Gaussian filter kernel and the depth buffer D. We specify the kernel size of G as percentage of the image diagonal, usually we use kernel sizes of 2% to 5%. Similarly to common local contrast enhancement, this function contains *local depth contrasts*, which can be interpreted as the following: $\Delta D \approx 0$ represents areas that are spatially not important, while $| \Delta D | > 0$ represents areas that are of special interest for our approach. Thereby, a negative spatial importance $\Delta D < 0$ represents areas of background objects that are close to other occluding objects, while a $\Delta D > 0$ represents boundary areas of foreground objects. As a result, spatially important edges, e.g. areas containing large depth differences, are located at singularities of this function. The obtained information is now used to alter only those areas that belong to objects of different depth.

(a) (b) (c) (d)

Figure 4: 3D example scene: (a) original rendering using a standard flat shading; (b) adding the spatial importance function ($\lambda < 0$); (c) linear combination of the original and a high contrast rendering; (d) depth darkening ($\lambda > 0$).

3.1 Image Modification Using the Spatial Importance Function ΔD

Following, we demonstrate several applications of the spatial importance function to modulate the color, luminance or contrast of the input image, thus, producing stronger depth cues for the recognition of the spatial arrangement. As already mentioned, the spatial importance function contains information about spatially relevant edges. This information can be integrated by directly modulating the input image I. Consequently, we add the spatial importance function to each color channel by computing

$$I' = I + \Delta D \cdot \lambda \quad (2)$$

with $\lambda \in \mathbb{R}$ being a user defined weighting parameter. The results usually provide additional depth cues, since the luminance of the input image is altered along spatially important edges (illustrated in Figure 3(c)). An important characteristic of this simple procedure is the interdependency of the depth differences and the amount of enhancement: the larger the depth differences, the higher is the visual enhancement. A drawback of this simple method is that we only consider depth differences. This can also reduce the visual contrast along an edge, if the color of the foreground object is brighter than the background (illustrated by the arrows in Figure 3(c)).

This is the reason for additionally considering the color of the input image (see Figure 3(d)). We preliminarily compute a version of the input image with an increased global contrast $C_{global}(I)$ using standard techniques known from common image editing tools, e.g. Histogram equalization or stretching. We can then use a linear combination of the original input image and the contrast enhanced version weighted by the absolute value of the spatial importance function ΔD:

$$I' = C_{global}(I) \cdot \mid \Delta D \mid \lambda + I \cdot (1 - \mid \Delta D \mid \lambda). \quad (3)$$

The contrast of the resulting image is locally increased depending on the spatial information that is obtained from ΔD. While this method enhances the visual contrast along spatially important edges, it requires a color or luminance contrast to be contained in the input image. Hence, this method fails if there is no color difference along an spatially important edge.

A possible solution for a missing color contrast is to introduce one kind of artificial contrast: similar to a cast shadow that is available in common image editing tools, we only darken the background along spatially important edges by computing

$$I' = I + \Delta D^{-} \cdot \lambda \quad (4)$$

with ΔD^{-} representing only the negative fraction of the spatial importance function. In this case only the more distant objects at a spatially important edge are darkened, while the foreground is left unchanged (illustrated in Figure 3(e)). The advantage is that we achieve convincing results that almost always emphasize the spatial relation between objects. As a consequence, the resulting image provides quite natural results that can be applied to almost any input data. This setting can also be utilized to enhance the spatially arrangement of a complex scenes. Here this technique slightly resembles an appearance that is known from global illumination techniques (see also Figure 5).

3.2 Further Examples

Analogous to the already described enhancement functions we can introduce various other effects that alter the input image. A number of further examples with their corresponding formal descriptions are described in the following.

Figure 4 shows a small 3D example scene with three rectangular objects one behind the other using a standard flat shading algorithm without shadows. This scene exemplifies a typical class of application for our algorithm: a white object in front of a white background, which cannot be visually detached due to their equal color. Figure 4(b) demonstrates the first case of image enhancement functions (2): the spatial importance function is directly integrated by adding it to the input image. As a result, we achieve an additional depth cue allowing us to visually separate all objects, but with decreased contrast at some borders, e.g. along the border of the white and the dark object. This problem is avoided in Figure 4(c) that demonstrates the second image enhancement function (3): the result is a linear combination of a high contrast version and the original image weighted by the absolute value of the spatial importance function. However, the white object cannot be separated from the background due to the missing luminance contrast. In Figure 4(d) the depth darkening (4) is applied introducing a very natural depth cue that, in our view, provides the best object separation.

Depth darkening is also especially useful in the case of many small isolated objects as it is usually found in complex botanical scenes. Figure 5 shows the result using a low-pass filtered depth buffer with a Gaussian kernel size of 2% of the image diagonal. With this technique the interior of the tree is darkened and the resulting contrast is increased. A similar effect is difficult to achieve using standard shading techniques, possibly by using precomputed lighting maps. In the accompanying video this effect can also be observed for moving objects.

A similar result can be achieved for a global illumination rendering of the Sponza Atrium scene as shown in Figure 6. Due to the global illumination the contrast is reduced in the shadowed regions in the scene. We are able to add contrast especially in spatially interesting regions such as under the arches. Of course, the differences must be more subtle in order to not influence the global illumination too much. In addition to the depth darkening in Figure 6(b), we applied depth brightening in Figure 6(c), which resembles a smooth halo effect. The latter effect is simply achieved by negating the user parameter λ.

Another possibility is to integrate color modulations: similar to the Gooch lighting model for technical illustrations [Gooch et al. 1998], we are able to integrate some artificial color contrast by modulating the foreground and the background with complementary colors, e.g. a cold and a warm color. Analogeous to the Gooch lighting model, we choose two colors, a cold one c_c and a warm one c_w. These colors are modulated by the negative fraction ΔD^- and the positive fraction ΔD^+ of the spatial importance function, respectively. Both factors are additionally weighted by a user parameter λ to specify the intensity of the effect:

$$I' = I + (\Delta D^- c_c + \Delta D^+ c_w) \cdot \lambda \tag{5}$$

Figure 7(c) demonstrates the effect: we add a shading that is inspired by the Gooch shading model in areas with large depth differences. As a result, yellowish halos appear around the objects that are enhanced by a blueish background shading.

The application of our functions to a photograph containing depth information is shown in Figure 8. Here we applied depth darkening using a Gaussian filter with a kernel size of 5% of the image diagonal. Especially for this example, which has a high color saturation, our effect provides additional depth cues that support the spatial recognition of the objects. The accompanying video demonstrates further photographic examples that are modified by our algorithm. Furthermore, we also show a small video animation that was captured with approximative depth information. Since this depth information is often reconstructed by image-based algorithms using stereo image pairs, visible artifacts are produced in smooth image areas that contain hardly any texture. However, the missing texture helps us to suppress these artifacts when applying our technique: similar to (3) that only produces additional depth cues if $\mid I - C_{global}(I) \mid > 0$, we compute a version of the input image $C_{local}(I)$ using a local contrast enhancement operator such as described by McHugh [2005]. As a result, $\mid I - C_{local}(I) \mid \approx 0$ in smooth and untextured image regions, which results in hardly any image enhancement and in turn suppresses visible depth artifacts. The result is demonstrated in the accompanying video, which shows additional depth cues with hardly any visible artifacts.

In Figure 9 we show the comparison of the result using ambient occlusions and our rendering. The effect is quiet different. While ambient occlusion simulates global illumination—and needs much precomputation time for complex objects—our method works in screen space, which allows for a better perception of the object details in a zoomed image. Both techniques can be combined.

3.3 Implementation

We implemented all operations as fragment shaders in order to exploit the computation power of graphics hardware, and to avoid the backreading of the depth information from the graphics board. This allows us to achieve reasonable high frame rates even for complex scenes. All the computer-generated images presented in this paper can be computed at a frame rate of 5-20 frames per second on a standard PC with 3 GHz and NVidia 6800 board. In general, the enhanced shading needs a fixed amount of GPU time for a given image size.

For a medium filter size of about 2% of the image diagonal the additional rendering effort is 40 ms for a 800×600 pixel image and 60 ms for a 1024×768 pixel image. If a larger filter kernel is used with a size of about 5% of the image diagonal our method takes 100 ms for a 800×600 pixel image and 150 ms for 1024×768 pixels.

4 Discussion

Since we consider differences between the original and the low-pass filtered depth buffer, the quality of the depth buffer is of special importance. This is particularly critical for photographs and videos that are captured including an approximated depth information. Here, artifacts may become visible that occur due to significant differences between the depth buffer and the color image.

Another issue is the non-linear distribution of depth values, meaning that there are significant amplitude variations of our spatial importance function when comparing objects in the foreground and objects in the background. However, this characteristic often provides natural results by emphasizing spatial relations in the foreground while hardly changing distant objects. It is also possible to recompute the depth values to achieve a linear distribution. Our examples were created using both non-linear and linear distributed depth values.

In addition, a significant attribute is the size of the filter kernel that is used to compute a low-pass filtered version of the depth buffer. A large filter kernel results in a larger area of influence, while small filter kernels reproduce more details, but at the same time tend to produce exaggerated results. Especially for thin objects there are significant differences: a large filter kernel smoothes out a thin object resulting in only minor image manipulations around that object, while a small filter kernel shows also here visible changes. Here the incorporation of additional parameters—similar to the threshold that is part of the conventional unsharp mask—might be helpful. This issue can be addressed in a future work looking for some adaptive approach.

An advantage of our approach is its general applicability. We can apply our algorithms to any image data with available depth information. Due to its simple and efficient realization using the capabilities of graphics hardware, we can utilize it in real-time graphics pipelines. We mainly use per-pixel operations, thus the computational effort mainly scales with the image resolution and depends only minorly on the scene complexity. Given coherent depth information, our method also provides a high frame-to-frame coherence during an animation due to its image filter characteristic.

In conclusion, the basis of our method is the spatial importance function that contains information about spatially important areas, which can be used in various ways to perform image manipulations. This great diversity sometimes makes it difficult to specify appropriate functions and parameters for a particular application.

5 Conclusion and Future Work

We presented a simple method for enhancing images that contain depth information. The difference between the original and the low-pass filtered depth buffer is computed in order to find spatially important areas. This information is utilized to locally enhance the contrast and the color of the input image, and thus, to improve the perceptual recognition. The method is useful for all scenes that contain spatially complex arrangements of objects. Applying our technique allows us to emphasize objects in the foreground and to visually depict the spatial relations in complex scenes. Here, especially the effect of depth darkening, meaning that we slightly darken background objects, introduces a natural additional depth cue by usually increasing the local contrast. Our method can be combined with all image operations that change the contrast, lightness, color or other parameters.

Our approach relies solely on color and depth buffer operations. Given that this data is coherent during an animation, our method provides also coherent results. Even for complex objects, e.g. botanical objects, stable results without artifacts are achieved as shown in the accompanying video. Furthermore, even in the case of editing video material containing depth buffer artifacts—which often appears if stereo vision approaches are used to reconstruct the depth buffer—we provided an enhancement operator that sill produces smooth results.

This work only gives a small selection of possible operators for depth dependent image manipulations. There are countless possilities for integrating the proposed spatial importance function into the image enhancement process. A possible direction in the future is to combine our method with non-photorealistic rendering techniques in the sense that the contrast enhancement is used to support additional hatching, stippling or other stroke-based techniques in order to increase the dynamics of such drawings.

6 Acknowledgments

We thank Daniel Scharstein and Richard Szeliski for providing the photographs with depth [Scharstein and Szeliski 2003] and the Visual Media Group at Microsoft Research for providing the video with depth information that is used in the accomplishing material [Zitnick et al. 2004]. The image of the Sponza Atrium model in Figure 6 was provided by Christoph Welkovits. Additionally, we want to thank the anonymous reviewers for their constructive suggestions that helped us to prepare the final version, and Anna Dowden-Williams for proofreading the paper. This work was supported by DFG Graduiertenkolleg (1042) 'Explorative Analysis and Visualization of Large Information Spaces' at the University of Konstanz, Germany.

References

ADAMS, A. 1980. *The Camera*. The Ansel Adams Photography Series. Littel, Brown and Company.

BEGHDADI, A., AND LE NEGRATE, A. 1989. Contrast enhancement technique based on local detection of edges. *Computer Vision, Graphics, and Image Processing 46*, 2, 162–174.

BUNNELL, M. 2005. Dynamic ambient occlusion and indirect lighting. In *GPU Gems 2*. Addison-Wesley, 223–233.

CIGNONI, P., SCOPIGNO, R., AND TARINI, M. 2005. A simple normal enhancement technique for interactive non-photorealistic renderings. *Computer & Graphics 29*, 1 (feb), 125–133.

DEUSSEN, O., AND STROTHOTTE, T. 2000. Computer-generated pen-and-ink illustration of trees. In *Proceedings of ACM SIGGRAPH 2000*, 13–18.

DICARLO, J. M., AND WANDELL, B. A. 2000. Rendering high dynamic range images. In *Proceedings of SPIE: Image Sensors*, vol. 3965, 392–401.

EYSENCK, M., AND KEANE, M. 2000. *Cognitive Psychology*. Psychology Press.

GOOCH, A., GOOCH, B., SHIRLEY, P., AND COHEN, E. 1998. A non-photorealistic lighting model for automatic technical illustration. In *Proceedings of ACM SIGGRAPH 98*, 447–452.

HUMMEL, R. A. 1975. Histogram modification techniques. *Computer Graphics and Image Processing 4*, 3 (sep), 209–224.

LEDDA, P., CHALMERS, A., TROSCIANKO, T., AND SEETZEN, H. 2005. Evaluation of tone mapping operators using a high dynamic range display. *ACM Transactions on Graphics 24*, 3 (aug), 640–648.

MCHUGH, S., 2005. Digital photography tutorials. http://www.cambridgeincolour.com/tutorials.htm.

MEYLAN, L., AND SÜSSTRUNK, S. 2004. Bio-inspired color image enhancement. In *Proceedings of SPIE: Human Vision and Electronic Imaging*, vol. 5292, 46–56.

NEYCENSSAC, F. 1993. Contrast enhancement using the laplacian-of-a-gaussian filter. *CVGIP: Graphical Models and Image Processing 55*, 6, 447–463.

PHARR, M., AND GREEN, S. 2004. Ambient occlusion. In *GPU Gems*. Addison-Wesley, 279–292.

RASKAR, R., TAN, K.-H., FERIS, R., YU, J., AND TURK, M. 2004. Non-photorealistic camera: Depth edge detection and stylized rendering using multi-flash imaging. *ACM Transactions on Graphics 23*, 3, 679–688.

REINHARD, E., AND DEVLIN, K. 2005. Dynamic range reduction inspired by photoreceptor physiology. *IEEE Transactions on Visualization and Computer Graphics 11*, 1 (jan), 13–24.

REINHARD, E., STARK, M., SHIRLEY, P., AND FERWERDA, J. 2002. Photographic tone reproduction for digital images. *ACM Transactions on Graphics 21*, 3 (jul), 267–276.

SAITO, T., AND TAKAHASHI, T. 1990. Comprehensive rendering of 3-d shapes. *Computer Graphics (Proceedings of ACM SIGGRAPH 90) 24*, 4, 197–206.

SCHARSTEIN, D., AND SZELISKI, R. 2003. High-accuracy stereo depth maps using structured light. In *Proceedings of Computer Vision and Pattern Recognition*, 195–202.

STARCK, J., MURTAGH, F., CANDES, E., AND DONOHO, D. 2003. Gray and color image contrast enhancement by the curvelet transform. *IEEE Transactions on Image Processing 12*, 6, 706–717.

STARK, J. 2000. Adaptive image contrast enhancement using generalizations of histogram equalization. *IEEE Transactions on Image Processing 9*, 5 (may), 889–896.

WINKENBACH, G., AND SALESIN, D. 1994. Computer-generated pen-and-ink illustration. In *Proceedings of ACM SIGGRAPH 94*, 91–100.

ZITNICK, C. L., KANG, S. B., UYTTENDAELE, M., WINDER, S., AND SZELISKI, R. 2004. High-quality video view interpolation using a layered representation. *ACM Transactions on Graphics 23*, 3 (aug), 600–608.

Figure 5: Enhancement of a complex botanical object using depth darkening.

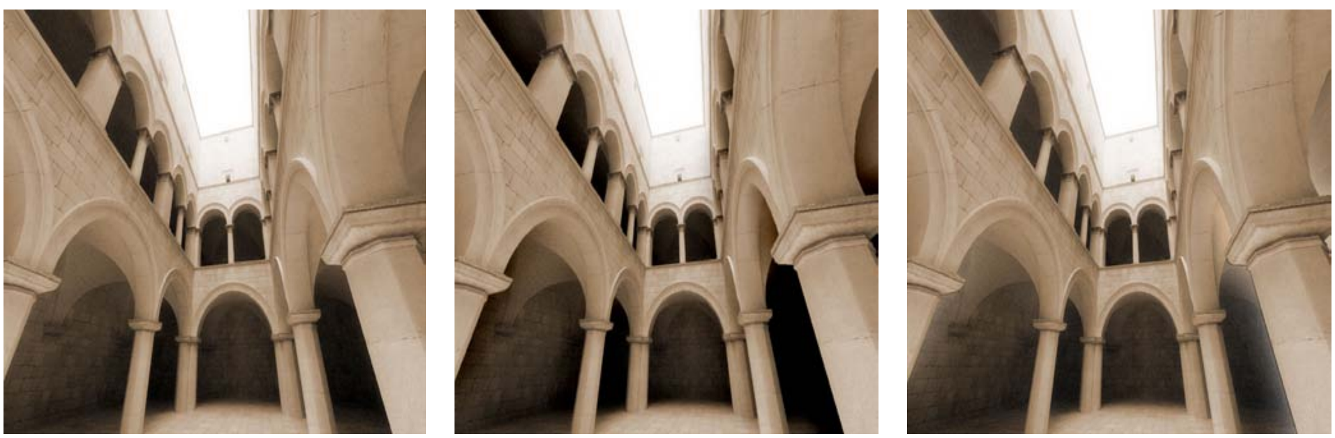

(a) (b) (c)

Figure 6: Example of improving a radiosity rendering of an architectural scene: (a) original rendering; (b) depth darkening; (c) depth brightening resulting in a halo effect.

Figure 7: Enhancement of a complex scene: original rendering (left); depth darkening (middle); colored version inspired by the Gooch lighting model (right).

(a) (b)

Figure 8: Enhancing a photograph with depth information (obtained from Scharstein and Szeliski [2003]): (a) original photograph; (b) depth darkening.

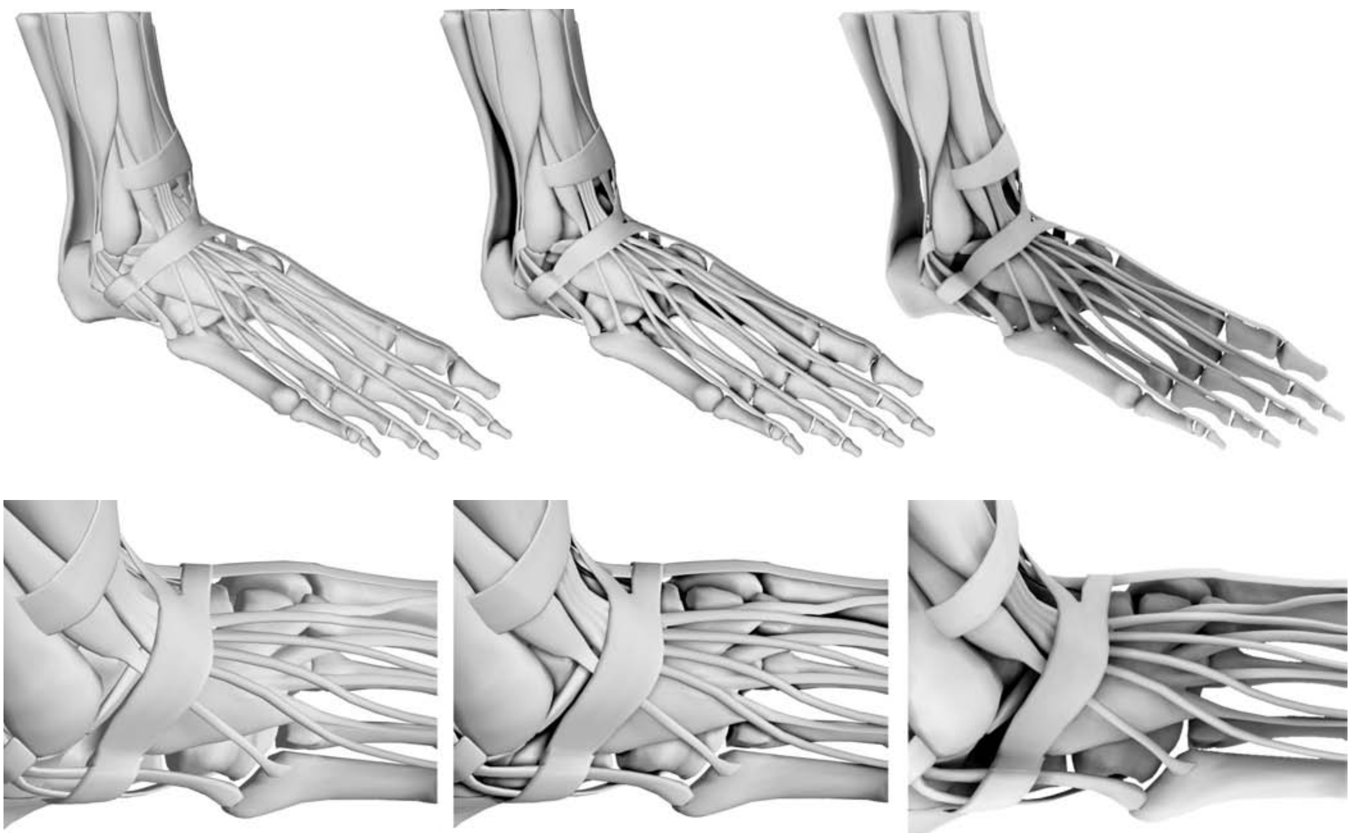

(a) (b) (c)

Figure 9: Comparison to ambient occlusion: (a) original rendering; (b) depth darkening; (c) ambient occlusion. While the ambient occlusion mimics a global illumination effect, our depth darkening operation works in screen space and allow for a better perception of details in a zoomed view.