Protecting identity documents with a just noticeable microstructure

Nicolas Rudaz, Roger D. Hersch
Ecole Polytechnique Fédérale
Lausanne, Switzerland

ABSTRACT

The development of plastic card printers has led to the widespread use of identity documents printed on plastic cards, such as ID cards, driving licenses and access key cards. This paper presents a new security feature based on a technique for embedding a personalized microstructure into a photograph. This microstructure takes the form of a bitmap pattern embedded into the original photograph as a succession of balanced chromatic shifts. The amplitude of these shifts may be modulated so as to make the pattern fully apparent, just noticeable, or completely invisible under normal viewing conditions. Since the chromatic shifts cancel each other out in any macroscopic portion of the image, the global appearance of the protected image remains intact. The embedded microstructure may be adapted to each instance of the protected identity document. For example, it can repeat textual information already present elsewhere on the document, or it can include a code derived from data specific to the document holder. Furthermore, this information may be made fully readable without the help of a specialized apparatus. Such identity documents exhibit an intrinsic resistance against imitation, tampering and substitution.

Keywords: security printing, identity document, halftoning, watermarking

1. INTRODUCTION

We are interested in creating color images and security documents with two layers of information, one layer of information being the global image and the second layer of information represented by a meaningful microstructure that is embedded within the global image. When seen from far away, only the global image should be visible. From a close viewing distance, the artistic microstructure should provide the main visual impact.

In the present contribution, we explore a new approach for incrusting human-readable microstructures into color images. The incrustation process yielding the final image embedding the two layers of information is automatic. Automatic microstructure incrustation paves the way to the mass production of individually customized security documents.

We incrust artistic microstructures into color images by synthesizing color differences between the foreground and background of the microstructure shapes. When viewed from a large distance, microstructure shapes embedded in the global image should not be distinguishable by the human eye. The resulting perceived color should be a combination of foreground and background colors, weighted by their respective surfaces. We assume that the color reproduction process offers enough resolution to create any desired color both for the foreground and the background shapes. We have therefore much freedom in creating color differences expressed in terms of lightness, saturation or hue difference components. This freedom represents a significant progress, compared with methods [1] that modify the lightness and preserve the chromaticity at locations specified by a mask (text message).

The first problem consists in creating continuous color differences within a 3D color space. A second problem is the creation of perceptually similar color differences across the color space. A further problem resides in determining the local coverage percentages of the microstructure's foreground and background, in order to appropriately choose their respective colors so as to ensure that when viewed from a certain distance, the combination of foreground and background appears as the desired target color.

EPFL/IC-LSP, CH-1015 Lausanne, Switzerland
Nicolas.Rudaz@epfl.ch
RD.Hersch@epfl.ch
http://diwww.epfl.ch/w3lsp/pub/publications/colour/

IS&T/SPIE 14th Ann. Symp. Electronic Imaging, Conf. Optical Security and Counterfeit Deterrence Techniques IV, 2002, SPIE Vol. 4677, 101-109.

2. CREATING THE MICROSTRUCTURE WITH COLOR DIFFERENCES

We would like to incorporate the microstructure into pictures by creating color differences that make the microstructure apparent. However, when viewed from far away, only the global image should remain visible. The microstructure can be made with repetitions of a symbol, a logo, text covering part of the image or ornaments. In the case of a text microstructure, a grid of characters is created and for each character, the respective relative foreground surface coverage is computed (Fig. 8f, color plate). The coverage values within the character's bounding-boxes remain constant. One may apply to the coverage map a smoothing filter in order to create smooth coverage values at the boundaries between character bounding-boxes.

To generate the resulting image, the output image is traversed pixel by pixel and corresponding positions within the original image and within the microtext layer are located. The color at the current location in the original image and the coverage percentage at the corresponding location in the microstructure (microtext) layer are determined. According to the coverage percentage, a color difference pair is formed and depending if the foreground or the background of the microstructure is located at the current location, the one or the other color of the color difference pair is selected as output color.

Fig. 8 illustrates this process with a portrait having the size of the photographs typically used in identity documents printed on plastic cards (Fig. 8a). A bitmap pattern composed by seamlessly tiling the string "COPYRIGHT/2001/EPFL/" is incrusted in this photograph (Fig. 8b). A 300% magnification of the bitmap pattern (Fig. 8c) and of the final image (Fig. 8d) is provided in order to better appreciate the color differences generated by the incrustation. The text is incrusted into the image with a color difference incorporating a small lightness and strong hue and saturation differences. The lightness and the saturation differences are apparent in the whole image. The hue difference is more subtle, but it can be easily observed under a strong magnification (Fig. 8e). At the same time, the original photograph and the processed image are nearly indistinguishable when seen from a certain distance.

One may generate color differences along the lightness axis, the saturation axis, the hue axis, or combined differences. A pair of difference colors forms in a linear 3D color space a vector pair having opposite directions. Their respective length ratio is inversely proportional to their corresponding relative surface coverage values (Fig. 9). The color working space must be linear, i.e. the sum of background and foreground colors weighted by their respective relative surface coverage values must generate the resulting target color. When the foreground surface vanishes, i.e. the microstructure disappears, the background color becomes the color of the original image. Color differences along the saturation axis are shown in Fig. 9b and along the hue axis in Fig. 9c.[1]

3. SYNTHESIZING COLOR DIFFERENCES

One would like to exploit the design freedom offered by color differences not only to create lightness variations, but also variations in saturation and hue. For the purpose of reasoning in terms of lightness, saturation and hue color differences, we apply to the RGB color cube a linear transformation comprising a rotation which places the black-white axis in vertical position.

On the planes perpendicular to the black-white axis, saturation and hue are represented in polar coordinates as radius and angle. The resulting coordinate system is named the LEF coordinate system [2]. The LEF lightness (black-white axis) is given by the L coordinate, and the LEF saturation-hue plane is given by the E coordinate (orientation: perpendicular to the L axis though Red) and the F-coordinate (perpendicular to the L and E axes) (see Fig. 1). The LEF coordinate system offers a good compromise between two divergent requirements: other color systems based on lightness (also called value), saturation (also called chroma) and hue such as HLS, HVC, CIELAB [3] are not linear and therefore they cannot be directly used for generating colors whose weighted sum should correspond to a given target color. The YIQ color system [4] is linear, but it is less intuitive than the proposed LEF system.

1.Due to simultaneous contrast, gray surrounded by blue appears as brown.

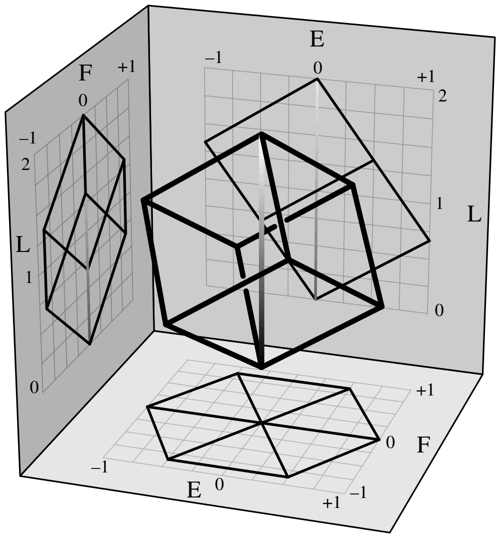

$$\begin{bmatrix} L \\ E \\ F \end{bmatrix} = \begin{bmatrix} \frac{2}{3} & \frac{2}{3} & \frac{2}{3} \\ 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & \frac{\sqrt{3}}{2} & -\frac{\sqrt{3}}{2} \end{bmatrix} \cdot \begin{bmatrix} R \\ G \\ B \end{bmatrix}$$

$$\begin{bmatrix} R \\ G \\ B \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & \frac{2}{3} & 0 \\ \frac{1}{2} & -\frac{1}{3} & \frac{1}{\sqrt{3}} \\ \frac{1}{2} & -\frac{1}{3} & -\frac{1}{\sqrt{3}} \end{bmatrix} \cdot \begin{bmatrix} L \\ E \\ F \end{bmatrix}$$

Fig. 1. The transformed RGB cube viewed in LEF space and the corresponding linear transformation

Fig. 2 shows constant lightness and constant hue families of slices in the LEF color space. Note that the intersection of a constant lightness slice and a constant hue slice is a radial segment starting at the black-white axis (achromatic color) and ending at the most saturated color located on the volume's external surface.

We cannot expect designers to manually specify color differences within all relevant parts of the color volume.[1] We therefore need a means for creating consistent color differences across the color space according to an initial specification of the amount of LEF lightness, saturation and hue differences. The principle we follow is to create for any desired color a color difference according to the initially specified color difference. We rotate the initial difference vector around the white-black axis and translate it in the radial direction towards more saturated colors until it becomes positioned at the desired location in the LEF color space. Our only aim at this stage is to ensure that for any wedge crossing a part of the color space, color difference vectors are continuous, i.e. no false contours are generated. In section 4 we introduce a method for rectifying color differences in order to produce perceptually similar color differences.

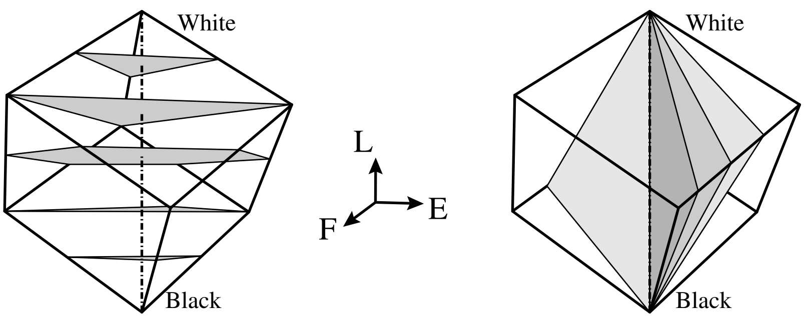

Fig. 2. Constant lightness and constant hue families of slices in LEF color space

Lightness differences can be created at any point inside the color volume[2]. However, since one achromatic color can be rendered by an infinite number of color pairs of opposite hues, it makes no sense to define saturation and hue differences on achromatic colors, i.e. on gray tones situated along the black-white axis. For this purpose and in order to ensure continuous color differences on a wedge crossing the black-white axis we require that achromatic colors have only an LEF

1. Relevant parts of the color volume are the parts containing colors present in the input image.
2. On points close to the volume boundaries, lightness differences have a range limited by the point's vertical distances to the volume boundaries.

lightness difference[1]. For colors located between the initially specified rotated color difference and the black-white axis, color differences are obtained by interpolating between the initially specified rotated color difference and the corresponding lightness difference. This interpolation transforms saturation and hue differences into lightness differences as a function of proximity to the black-white axis.

At the boundary or close to the boundaries of the LEF volume, it is not possible to preserve the initial color difference. We adopt a best effort strategy by modifying when necessary the original saturation difference and keeping as much as possible the hue difference. Fig. 3 shows how saturation and hue differences are adapted in order to remain within the boundaries of the color volume.

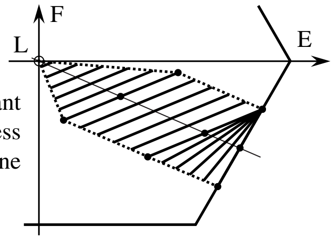

Fig. 3. Adapting saturation and hue differences to remain within the boundaries of the color volume

If necessary, the initial lightness difference is reduced so as to remain within the boundaries of the color volume. As shown in Fig. 4, at positions close to the volume's boundary, we interpolate between the position where the color difference vector touches the volume's boundary and the position where the color difference vector is projected onto the volume boundary and resized. Resizing is carried out in order to avoid an enlargement of the absolute color difference.

At the color volume's vertices, all color differences are reduced to zero, i.e. no barycentric combination of colors exists which can render the basic colors Black, White, Red, Green, Blue, Cyan, Magenta, and Yellow.

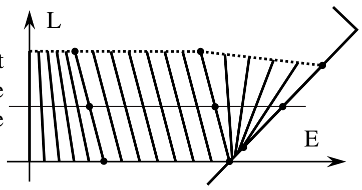

Fig. 4. Adapting LEF lightness and saturation differences at the boundary of the color volume

Such vanishing color differences are tolerable and may even produce interesting effects. If color differences need to be avoided, one can slightly desaturate the input image and ensure that it does not incorporate colors located at the color volume's vertices. Then it will be possible to produce color differences within all image parts.

4. PERCEPTUALLY SIMILAR COLOR DIFFERENCES

Our aim is to produce perceptually similar color differences across all or at least most colors of the color volume. It is known that equal color differences (in terms of Euclidean distance) in RGB or CIE-XYZ space do not produce perceptually equal color differences. It is however less known that equal color differences in CIELAB, i.e. color differences with the same significant ΔE value (e.g. CIELAB $\Delta E = 10$) do not produce perceptually equal color differences either.

In certain regions of the LEF space, the perceived color differences become rather low. A more detailed analysis of this phenomenon reveals that this occurs in regions where there is an inversion in the CIELAB lightness differences, i.e. where CIELAB lightness differences become zero and then grow in the opposite direction. Fig. 5 shows constant light-

1.The module of the lightness difference can be made equal to the module of the initially specified color difference.

ness curves (fixed L in CIELAB) of a given hue slice in the LEF space. The literature mentions that for small color differences (just noticeable differences), lightness differences and chromatic differences seem to be additive. However when the lightness difference is sufficiently high, it masks the contribution of the chromatic color difference [5].

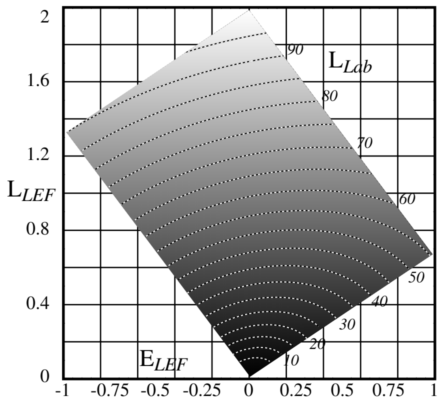

Fig. 5. Constant CIELAB lightness curves within a hue slice in LEF space

In order to create color differences as perceptually similar as possible across the whole color volume, and to avoid CIELAB lightness inversions when producing LEF hue and saturation differences, we constrain our technique to use LEF color differences incorporating at least a small CIELAB lightness difference (i.e. a lightness difference of $\Delta L_{Lab} \geq 2$).[1] After having synthesized an LEF color difference according to section 3, a last lightness difference correction step is carried out in order to preserve the originally specified CIELAB lightness difference. For this purpose, we convert the generated LEF difference to CIELAB and compute by what percentage the LEF difference needs to be reduced or enlarged to fit the CIELAB ΔL lightness difference implicitly specified in the initial color difference. Several iterations are needed to converge to the final corrected LEF color difference. These iterations produce a corrected LEF color difference incorporating a CIELAB lightness difference ΔL within ±0.2 of the initially specified lightness difference.

5. COVERAGE ESTIMATION

When seen from a certain distance, the foreground color C_f and background color C_b of the microstructure contribute to the resulting perceived color C according to their respective surface coverages α and $(1 - \alpha)$:

$$C = \alpha\, C_f + (1 - \alpha)\, C_b \quad ; \quad 0 \leq \alpha \leq 1$$

Since we would like to compute for a given target color C the respective colors of foreground C_f and background C_b, it is important to deduce from the microstructure pattern the foreground and background coverages within a local neighborhood.

In the case of typographic characters, surface coverages are easily computed by considering the bounding-box of each character and determining the ratio between the discrete character surface (foreground) and the bounding-box surface (Fig. 8f). However, in the case of irregular structures of different sizes the coverages are more difficult to estimate. The

1.In order to adjust CIELAB lightness differences, one needs to transform color coordinates from CIELAB to the display RGB color coordinate system. Conversions from CIE-XYZ to CIELAB and vice-versa are defined in [6]. The conversion from the display RGB values to CIE-XYZ and vice-versa requires a calibrated display with known gamma values and known CIE-XYZ values of the red, green and blue color channels [7]. In our work, images incorporating color differences were generated on a LaCie electron21/108 calibrated display with a gamma value of 1.8 and printed on a dye sublimation Kodak 8670 PS printer using Kodak's color management software.

distance between neighboring structures gives some clue about the succession of foreground and background intervals as well as about the bounding-boxes within which the coverages percentages may be computed.

We explore the distance between neighboring structures by starting at the center of the microstructure and launching 16 lines at equally spaced angles, i.e. angles spaced apart by $2\pi/32$ (11.25 degrees). For each line (Fig. 6 left), we obtain an approximated period M_1M_2 of the microstructure comprising a segment ($T_{11}T_{21}$) of the starting color (e.g. foreground) and two half segments of the other color ($T_{11}M_1$, $T_{21}M_2$). We also obtain the local surface coverage C_R along that line. In order to compute a bounding-box corresponding approximately to one structure period, we compute for each main orientation (top, left, bottom, right) the mean location of the period end points (Fig. 6 right).

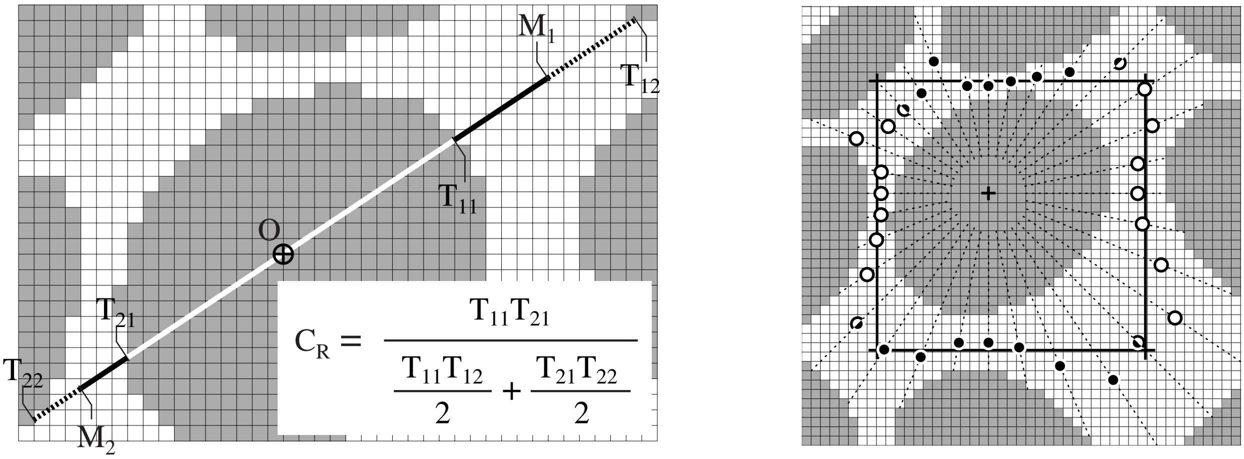

Fig. 6. Left: approximation of the local microstructure period. Right: determining the structure bounding-box by averaging in each main orientation (top, left, bottom, right) the location of the period end points.

We define the final resulting microstructure surface coverage for the analyzed local area as the mean between the surface coverage computed by taking the mean coverages of all lines and the surface coverage computed within the bounding-box. When carried out on synthetic examples, the computed microstructure surface coverage is close to the actual surface coverage (within 10%).

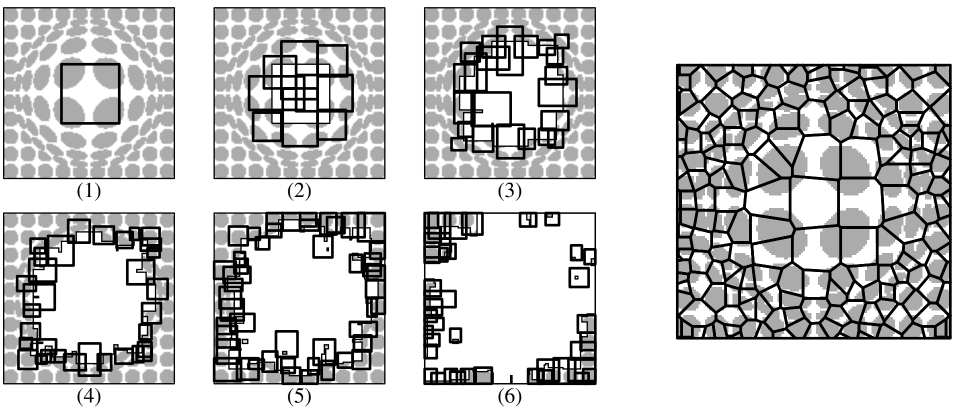

Fig. 7. Left: computation of coverages in successive regions of the microstructure. Right: Voronoi tesselation of the seed points.

The so computed bounding-box delimits the region for which the coverage is computed. Its vertices as well as its mid-points between vertices are the new seed points, from which new coverage computations according to the same algorithm are carried out. Seed points located within already computed coverage bounding-boxes are removed. Fig. 7 left shows the succession of computed regions within the microstructure. At the end of these coverage computation iterations, a local coverage value is associate to each valid seed point. The spatial extent of each seed point is determined by a Voronoi tesselation of the seed points (Fig. 7 right). The resulting polygons surrounding the seed points indicate the sub-spaces on which computed local surface coverages remain valid.

The coverage estimation algorithm described above adapts itself to the local period of the microstructure and yields good results in the general case, i.e. for bitmap patterns made of randomly distributed, variable-size arbitrary shapes.

6. SECURITY IMAGES PRODUCED BY COLOR DIFFERENCES

Incrustation of a microstructure into a color image can also serve as a visible watermark. In contrast to previous visible watermarking techniques [8] which modify the lightness of image parts where the watermark is present, microstructures encrusted with the help of color differences only create local color differences compensating one another. Due to the low pass filter behavior of the human eye, from a certain distance, only the original image with its original color appears.

In the photographs for identity cards shown in Fig. 8 and 10, the microstructure encrusted in the images is composed of characters and numbers. A verification code can be incorporated into the text or the numbers encrusted in the image. This text can undergo a geometrical transform, and it is also possible to make it run across the whole document in order to further improve its resistance against falsification (Fig. 11). The encrusted color differences are subtle and do not disturb the global face appearance. With a magnifying glass and under a good light source, one can easily recover the underlying text microstructure. The color differences are consistent and look uniform across a broad range of skin colors. Furthermore, the method works equally well on black-and-white photographs. Different variants of our method have been tested in collaboration with company Orell Füssli Security Printing Ltd, and the results we obtained confirmed indeed our expectations. A concrete first application for a driver license was developed and is in its testing phase now.

On a PowerPC G4 450 MHz computer, the processing time required for microstructure incrustation into one photograph (295 x 372 pixels) is inferior to 4 seconds. Incrustation time is proportional to the size of the final image.

7. CONCLUSIONS

The creation of artistic color images incorporating two levels of information, one at the global image level and one at the microstructure level is a challenging task. The microstructure incrustation tool we propose allows to create such images with a minimal amount of technical effort. Designers can concentrate on aesthetic aspects, i.e. choosing a suitable combination of global image and microstructure, as well as a masking layer specifying on which parts of the global image the microstructure is to be applied. After specifying an initial color difference, the incrustation software automatically embeds the microstructure into the global image by synthesizing appropriate color differences.

Generating faithful global images made of repetitive or non-repetitive microstructures required to elaborate solutions to two problems. The first problem is the synthesis of continuous color differences across the color space. We offer designers the design freedom of specifying color differences in lightness only, in lightness and saturation, in lightness and hue or combining lightness, saturation and hue. The initial color difference is propagated throughout the color volume. At the proximity of the white-black axis, it progressively transforms itself into a lightness difference. When it comes close to the volume boundaries, it progressively adapts itself to fit the color volume boundary surface.

The second problem is the computation of the local foreground or background coverages for a possibly non-repetitive microstructure incorporating significant variations in both coverages and local pattern periods. The robust solution we propose consists in specifying a seed point, throwing from that seed point lines at a large number of discrete orientations and analyzing the local pattern period along these lines. The mean values of the local pattern periods allow to create a bounding-box surrounding the local pattern. The final local coverage is a function of the line coverages and of the coverage within the bounding-box. Points on the bounding-box are used as new seed points for computing neighboring foreground and background coverages and for propagating the coverage computation iteratively to the full image.

The tool we created allows without much technical efforts to create security documents incorporating freely designed or scanned bi-level microstructures. The incrustation of characters and numbers into valuable document images (e.g. photographs in ID cards) provides a protection against counterfeiting attempts. The incrustation of such information can be specifically adapted to each instance of a security document, increasing its resistance against falsification by substitution.

ACKNOWLEDGMENTS

We would like to thank Orell Füssli Security Printing Ltd, especially MM. Coleman and Eichenberger for having contributed to this collaborative effort. We also thank the Swiss CTI fund for having partly financed the present research (contract CTI 3776-1).

REFERENCES

1. US patent 5,530, 759, "Color Correct Digital Watermarking of Images", inventors: W. Braudaway, K.A. Magerlein and F.C. Mintzer, issued June 25, 1996.
2. N. Rudaz, R. D. Hersch, V. Ostromoukhov, "Specifying color differences in a linear color space (LEF)", *Proceedings of the IS&T/SID 97 Color Imaging Conference: Color Science, Systems and Applications*, Nov 17-20, 1997, Scottsdale, Arizona, USA, 197-202.
3. J. Foley, A. van Dam, S. Feiner, J.F. Hughes, *Computer Graphics. Principles and Practice*, 2nd edition, Addison Wesley, 1990.
4. W. K. Pratt, *Digital Image Processing*, 2nd ed., John Wiley & Sons, New York, 1991.
5. T.L. Chen, C.Y. Yu, "The Relationship between Visual Acuity and Color Contrast in the OSA Uniform Color Space", *Color Research and Applications*, Vol 21, No. 1, Feb. 1996, 18-25.
6. G. Wyszecki, W.S. Stiles, *Color Science*, J. Wiley, 1982.
7. H. R. Kang, *Color Technology for Electronic Imaging Devices*, SPIE Optical Eng. Press, 1996, Chapter 14.
8. G. Braudaway, K. Magerlein, and F. Mintzer, "Protecting publicly available images with a visible image watermark", *Proc. SPIE: Optical Security and Counterfeit Deterrence Techniques*, Vol 2659, 1996, 126-133.

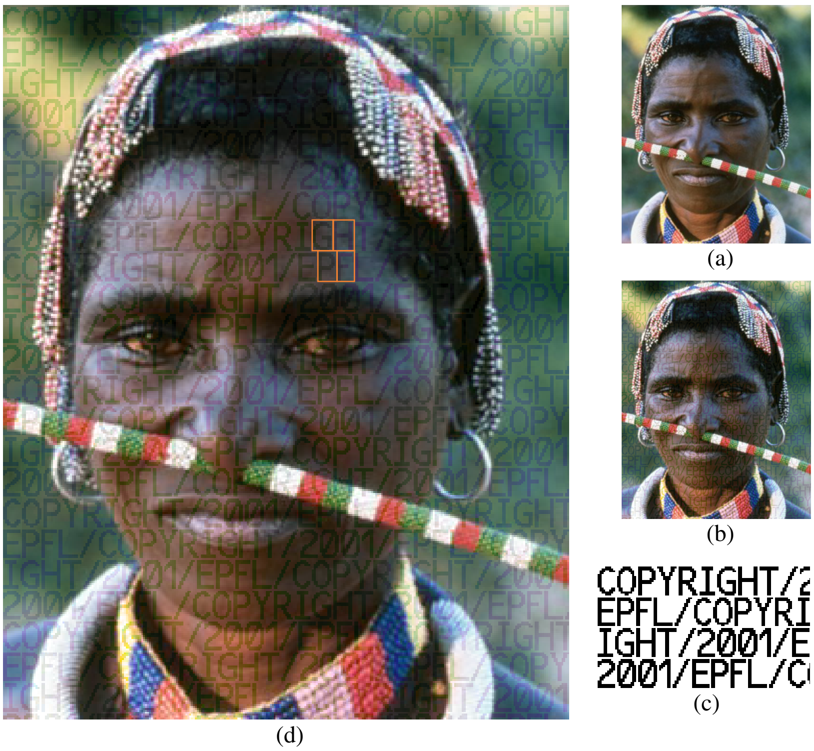

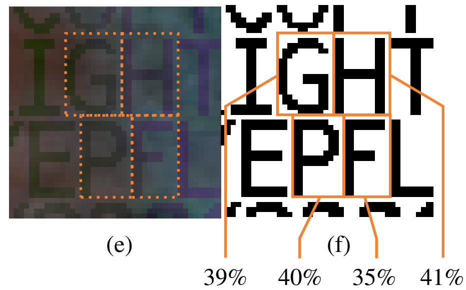

Fig. 8
Incorporating the microstructure into the image

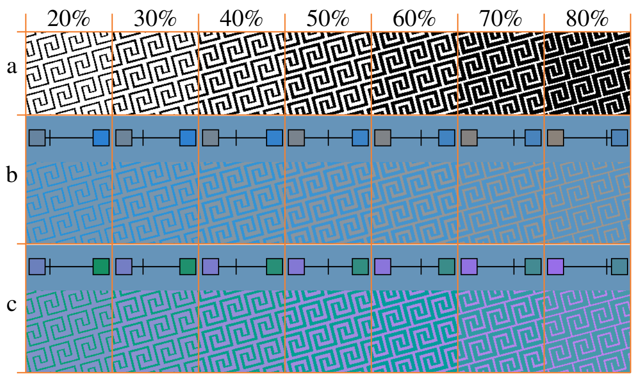

Fig. 9 Color differences for diffferent surface coverages

Fig. 10
Encrusted text as a visible watermark in monochrome and color photographs

Fig. 11
ID card with geometrically transformed text running over the whole document