# Skin-Sensitive Automatic Color Correction

Hila Nachlieli, Ruth Bergman, Darryl Greig, Carl Staelin, Boris Oicherman, Gitit Ruckenstein, Doron Shaked

HP Laboratories
HPL- 2009-13

**Keyword(s):**
skin detection, face detection, adaptive enhancement, object dependent enhancement, image probability map, color correction, sharpening

**Abstract:**
One of the most important aspects of a good image is that the subject looks good. It is therefore our objective, when building an automatic image enhancement system, to optimize the appearance of the image subject(s). Moreover, people often have specific preferences regarding the look of different objects. Grass, for example, is considered prettier when green and with sharp texture, while skin is expected to be smooth, and, preferably, not green. An image enhancement process that involves such object dependent considerations requires the guidance of a human supervisor able to identify image content and adapt the enhancement parameters for each content object accordingly.

We propose to automate content sensitive image enhancement to correspond to human preferences regarding the appearance of people in an image. We do this by detecting the skin tones of specific people in a given image. Based on these learned skin tones, we generate a fuzzy object map representing skin areas in the image. This map is used to automatically tune parameters of the enhancement algorithms. We present an adaptive skin color correction algorithm. We demonstrate the effectiveness of our algorithms in both a laboratory setting and in a production system.

External Posting Date: February 6, 2009 [Fulltext]
Internal Posting Date: February 6, 2009 [Fulltext]

Approved for External Publication

Submitted to SIGGRAPH 2009, New Orleans

# Skin-Sensitive Automatic Color Correction

Hila Nachlieli, Ruth Bergman, Darryl Greig[††], Carl Staelin, Boris Oicherman,
Gitit Ruckenstein[*], Doron Shaked
[†]HP Laboratories Israel, Haifa, Israel, [††]Multimedia Interaction and Understanding Lab.

**Abstract**

One of the most important aspects of a good image is that the subject looks good. It is therefore our objective, when building an automatic image enhancement system, to optimize the appearance of the image subject(s). Moreover, people often have specific preferences regarding the look of different objects. Grass, for example, is considered prettier when green and with sharp texture, while skin is expected to be smooth, and, preferably, not green. An image enhancement process that involves such object dependent considerations requires the guidance of a human supervisor able to identify image content and adapt the enhancement parameters for each content object accordingly.

We propose to automate content sensitive image enhancement to correspond to human preferences regarding the appearance of people in an image. We do this by detecting the skin tones of specific people in a given image. Based on these learned skin tones, we generate a fuzzy object map representing skin areas in the image. This map is used to automatically tune parameters of the enhancement algorithms. We present an adaptive skin color correction algorithm. We demonstrate the effectiveness of our algorithms in both a laboratory setting and in a production system.

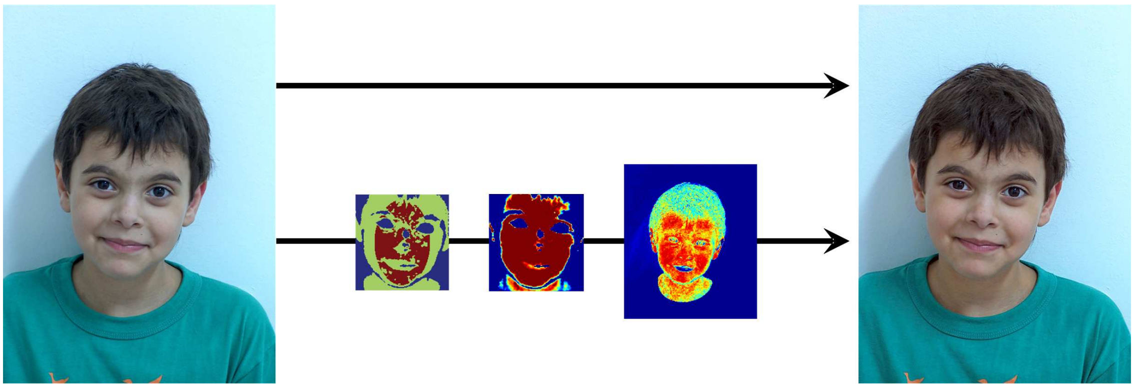

**Figure 1:** *Skin detection and color correction.*

## 1 Introduction

People have definite preferences and expectations regarding the appearance of images. They want images to be crisp and colorful, with clear details. They want the people they love to look good in the images. As digital printing becomes more accessible, more time and energy is expended preparing photo albums and

[*]Gitit Ruckenstein, our esteemed colleague and friend, passed away in August 2007
[†]E-mail: {hila.nachlieli,ruth.bergman, darryl.greig, carl.staelin, boris.oicherman, doron.shaked}@hp.com

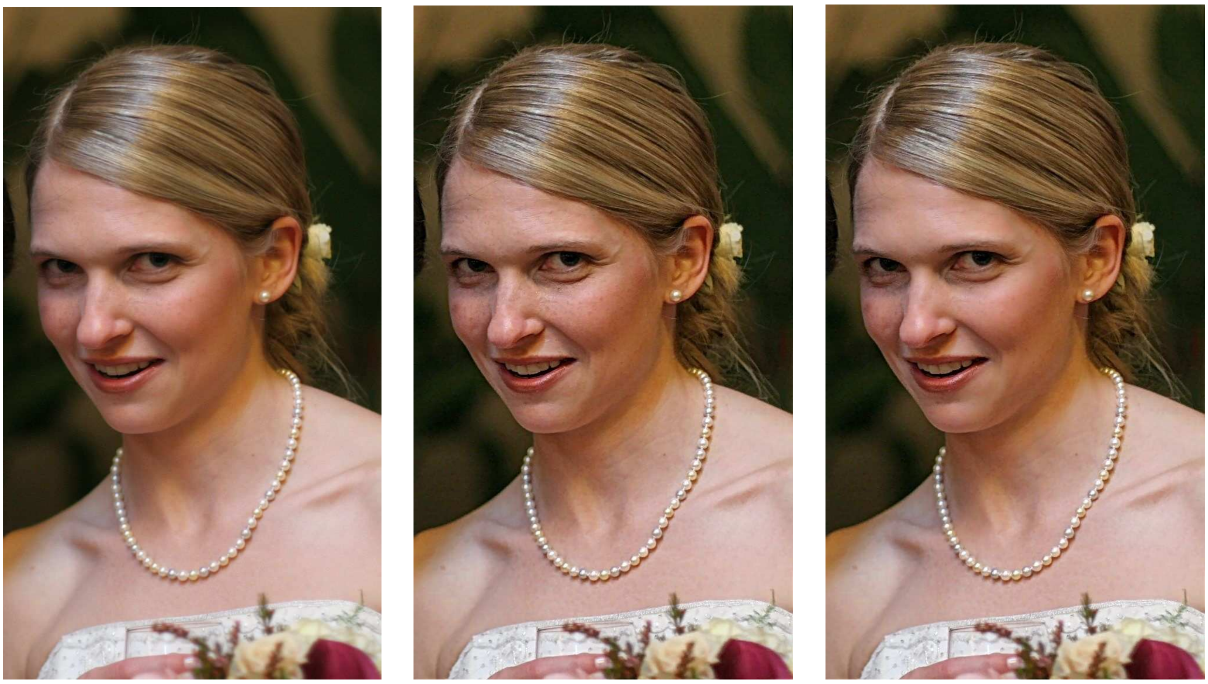

(a) (b) (c)

**Figure 2:** *Automatic skin-sensitive sharpening.(a) is the original image, (b) is global sharpening, where the lace and necklace, but also pimples and wrinkles are enhanced, and (c) is the proposed skin sensitive sharpening, where non-skin details such as eyes, lace and necklace are enhanced, yet the sharpening in skin regions is reduces, to avoid skin defect enhancement. The original image is taken from flickr, Copyright Notice : Creative Commons 2.0 / Attribution, http://www.flickr.com/photos/markjhandel/.*

calenders, and users expect the final product to meet their quality expectations. It is therefore desirable to fully automate the image enhancement process. Moreover, since one bad image can cause the customer to return an entire album, automatic image enhancement should improve, or at-least not damage, all of the processed images. However it is well known that globally applied enhancement techniques can improve some images but damage others. Automatic image enhancement should, therefore, enhance each region in each image by exactly the 'right' amount based on the attributes and content of that specific region. While available image processing algorithms use specific image analysis tools that respond to smooth areas, image edges and noise to guide the enhancement parameters, they lack the ability to enhance faces in the image according to human preferences. In this paper we propose to use automatic skin detection in order to automatically tune the parameters of image enhancement algorithms, thereby achieving automatic skin-sensitive image enhancement.

People, and in particular faces, are often the primary subject of an image. Attention is usually drawn to the faces, and people tend to have definite opinions regarding how a face should appear. This is particularly true if the face in the photo is of someone they know - not to mention their own. Viewers prefer facial details such as eyes, lips and hairline to be crisp, whereas skin details such as wrinkles and blemishes should be blurry and less visible. Healthy skin colors are naturally preferred over sickly looking ones. Therefore our objective is to automatically improve the appearance of faces in images. Clearly this requires automatically identifying faces in the image, however face detection alone is not enough, since our expectations of skin appearance differ significantly from our expectations regarding the appearance of other facial features such as eyes. Therefore an accurate skin map is required. Such a map should identify, with high probability, all skin pixels, and only skin related pixels, since some image enhancement algorithms (for example color correction) may create strange artifacts when applied to only portions of the skin regions, or when applied to non-skin objects.

Several attempts have been made to improve the appearances of faces in images. Most of these relay on user guidance, e.g. [1, 2, 3]. There has also been some initial work around automating object sensitive enhancement algorithms. One example is Skin-Aware Local Contrast Enhancement [4]. Since automated skin-aware image enhancement for real-world applications should be fool-proof, it requires accurate skin detection. The skin detection algorithm used in [4] is based on a global Gaussian skin color model. This approach is similar to our initial global skin color model (Section 2.2). It turns out that such a global skin model, by itself, does not provide the reliable guidance essential for automatic image enhancement. For example, it tends to incorrectly identify sand, autumn trees, and dry foliage as skin, and fails to detect less-conventional skin colors or skin which was captured under unusual luminance conditions.

In this article we present a skin detection algorithm which is sufficiently accurate to enable automatic skin-sensitive image enhancement. The skin detection algorithm is adaptive to the skin colors of the individual people in images. Figure 1 depicts the major steps of the adaptive skin detection algorithm. The algorithm uses face detection and a global skin color model to coarsely discriminate skin from non-skin pixels on the face (second left). For each face a learning procedure refines this skin labelling (middle) and the parameters of the skin color model are adjusted according to the color of the skin pixels. The per-face color models are used to compute the final skin probability map (second right).

The resulting skin map may be used to automatically tune several image enhancement algorithms. This paper describes two skin-sensitive image enhancement algorithms. Skin-sensitive automatic image sharpening is demonstrated in Figure 2 and further described in section 3.2. Skin-sensitive automatic color correction is demonstrated in Figure 1. The subject in the left image looks ill, whereas in the color-corrected image on the right her appearance is much improved. Color correction is presented in detail since this application best exemplifies the proposed approach.

Skin tones, along with blue sky and green foliage, belong to a group of colors termed *memory colors*. For each of the memory colors, people have a rather well defined *prototypes* stored in their mind. When looking at a picture, the viewers seem to test the 'match' of the image memory colors to that prototype. Skin that does not match the prototype often renders the whole image unpleasant. Therefore skin tones have received special treatment since the early days of color photography [5].

The *preferred* reproduction colors do not necessarily represent the real colors of objects. Thus, preferred skin tones are yellowish and more saturated [5, 6]. Cultural differences in color reproduction preferences were reported in [7, 8] although they are not expected to have significant practical effect [8]. It was also suggested that the "naturalness" of an image is determined by the appearance of its most important objects. Indeed, professional image editors tend to segment images into areas they consider most important, specifically skin, and handle the color enhancement in these areas independently [9]. The high precision of the skin maps obtained in such manual segmentation is essential to avoid maltreatment of other skin colored objects, and to prevent artifacts resulting from color correcting only part of the skin regions in the image. The proposed fully automatic algorithm was shown to either improve or not damage images when tested on thousands of images.

This skin-sensitive automatic color correction algorithm, along with the skin-sensitive automatic sharpening, is implemented in the HP-Indigo photo Enhancement Server (HIPIE), described in [10]. HIPIE is a commercial printing application that processes millions of images per day prior to printing without diminishing the quality of any of them. The fully automatic feature of these algorithms is also desirable for other real-life applications. These include video conferencing applications, where participants' appearance can be improved without degrading other objects in the image, such as items presented to the other participants, and High Definition TV (HDTV), where over-sharpening of faces is a known problem.

# 2 Face and Skin Analysis

The skin map describes the approximate likelihood that a given pixel corresponds to skin tone. The proposed fuzzy skin map is specific to an individual in an image, hence, if there is more than one person in an image we have a different skin map for each person. This adaptation to a specific person is introduced by learning the characteristic colors of the skin regions in the faces detected in the image.

Given an image, we start by calculating a global skin map, and independently locating the faces in the image. Next, we examine the faces one by one and use the global skin map information, as well as location information, to extract a fuzzy face-skin map for each face. We then learn the parameters of the skin color model that best describes the colors of the skin in that particular face, and create a body-skin map for the corresponding person in the current image. For some applications, e.g. adaptive sharpening, we combine the body skin maps to generate a unified per-image skin map.

## 2.1 Face detection

The face detection we use is a multi-view extension of the single (portrait) view face detector described by Viola and Jones [11]. The original Viola-Jones detector consists of low-level features measuring intensity differences between rectangular regions of the image. These features are applied at different scales and locations in a square candidate window to form a pool of weak classifiers (i.e. classifiers with only weak discriminating power between faces and non-faces). Figure 3 shows an example set of such features, and how a single feature applied in a square window might be expected to give some discriminating power between faces and non-faces.

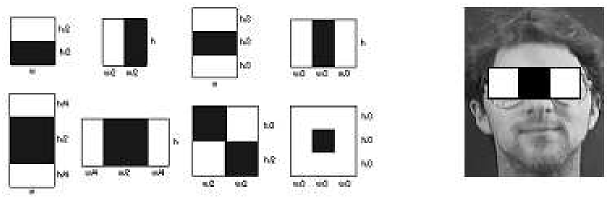

**Figure 3:** *Low-level face detection features and example application to a candidate patch.*

The features are combined into a sequence of stages using a variant of the Adaboost fitting method as described in [11]. Each stage is defined by a very high acceptance rate of true face examples (say 99%), but only a moderate rejection rate of non-face examples (say 50%). When a candidate image patch (x-y location and scale) is passed to a cascade it is processed by each stage in turn and either accepted, in which case it is passed on to the next stage, or rejected, in which case processing ceases. The face / non-face acceptance characteristics of the stages mean that most non-faces are eliminated early, while most faces are retained.

For example, given the rates above, after only six stages 98% of all non-face patches have been eliminated while 94% of all face patches remain. The result is that the detectors run very fast because the vast majority of image patches are non-faces, and most of these are eliminated early in the cascade processing.

The multi-view detector consists of a parallel set of cascades, each of which corresponds to a different face view (such as profile, inverted, $30^o$ rotated etc), and a patch that passes successfully through every stage of a cascade is labelled as a face in that cascade's view. During normal operation, the detector evaluates all the cascades in parallel and returns either a not-face verdict, or the view resulting from a successful cascade evaluation.

There are some practical restrictions arising from using a face detector in this way. In particular, the detector evaluation is an image-wise process, and cannot be implemented row-wise or strip-wise. Furthermore, even a very good detector will miss a significant percentage (5-10%) of obvious faces. This is mitigated somewhat by modifying the detection to give a confidence measure for each face (rather than a hard threshold), however the final usage still needs to be robust to the possiblity of both false positives and false negatives.

### 2.2 Global Skin Map

The global skin map computation is done for every pixel separately and is based only on the values of the pixel. In [12], Martinez studies a wide range of images that include people with different skin tones and in different luminance conditions. Martinez segments the skin regions out of every image and collects the R,G,B pixel values in each segment. He observes that the collected data is mapped to a compact blob of values in the LCH color space. He reports the mean and standard deviation values of that blob in each of the LCH coordinates. We use the findings of [12] to define a Gaussian probability function for skin

$$P(skin|l,c,h) = Z\ e^{-\left(\frac{(l-\mu_l)^2}{2\sigma_l^2}+\frac{(c-\mu_c)^2}{2\sigma_c^2}+\frac{(h-\mu_h)^2}{2\sigma_h^2}\right)} \tag{1}$$

where $(l,c,h)$ are the LCH coordinates of this specific pixel, $\mu_l = 181$, $\mu_c = 32$, $\mu_h = 34$, $\sigma_l = 30$, $\sigma_c = 11$, $\sigma_h = 8$, and Z is a normalization factor. The resulting function attaches high probabilities to a large spectrum of skin tones, while non-skin features typically attain lower probabilities. An example of our fuzzy skin map is shown in Figure 4(d), where skin has high likelihood values while the garment and background have low skin likelihood. On faces, the skin map distinguishes non-skin features, specifically eyes, from skin. The computation of the skin map involves a transformation of pixel values from RGB to LCH, a Gaussian probability computation in each of the LCH channels and a range adaptation function. The probability computation may be replaced with a look-up table that maps LCH color space values to skin probabilities.

### 2.3 Personal Skin Map

At this stage we have already calculated the global skin probability model (Section 2.2), and detected the faces in the image (Section 2.1). For each detected face, we crop a region around the face, and focus on this region for face skin detection. We resize the cropped image to a fixed scale and smooth it to reduce noise. An example is shown in Figure 4(c), which is the cropped image of the boy from 4(a).

We rely on each of the face boxes to contain face, and, hence, contain skin regions as well. This way even if the skin color is not a typical skin color, we know it is skin because it is part of the face. However, not all of the details in the face-box are skin. We assume that at least a quarter of the pixels in the face box are

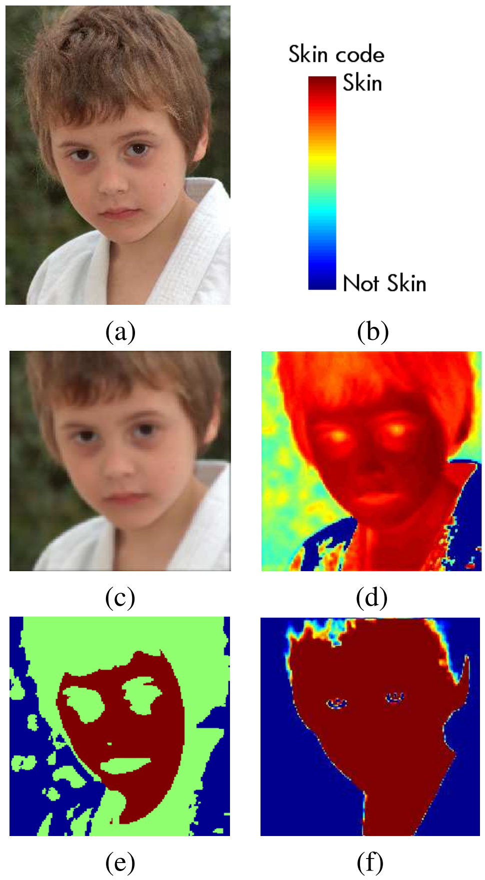


**Figure 4:** *Refining skin detection in face regions. (a) original image, (c) cropped image, based on face detection. (d) original skin map, based on global color-based model. Brown pixels are skin, blue pixels are non-skin, as shown by the color bar (b). (e) initial tags, given according to skin likelihood (d). Brown pixels are tagged as skin and blue pixels are tagged as non-skin. (f) final face-skin map.*

skin, that the skin pixels are located in the middle of the face box, and that the skin regions in the face have high skin likelihood values, at least relative to non skin regions in the face box. We use these assumptions to tag a quarter of the face box pixels as skin pixels. Non skin regions, such as eyes and background can be anywhere in the face-box, and we tag a quarter of the pixels, which includes the pixels having the lowest skin probability, as "non skin". The remaining pixels are marked as "do not know". This tagging is demonstrated in Figure 4(e), where the "skin" tagged pixels are marked as brown, "non-skin" tagged pixels marked as blue, and the "do not know" tagged pixel are green.

We find that five of the features of the pixels correlate with skin likelihood: location relative to the middle of the face box $(x, y)$ and color $(l, c, h)$. The importance of each of those features in separating skin from non skin varies between images and faces. For example, if the face is located in the left size of the face box, the x feature is highly important, while the y feature might be insignificant. We learn the relative importance, denoted by $w_l$, $w_c$, $w_h$, $w_x$, $w_y$, of each of those features adaptively, for each face box, by comparing the features of the pixels tagged as "skin" to the features of the pixels tagged as "not skin" and "do not know".

We experimented with two learning methods in order to identify the skin region inside the face box. The first approach is K-nearest neighbors, in which each of the "do not know" pixels gets face-skin likelihood based in a weighted average over it's 1000 nearest neighbors in the five dimensional feature space, which is defined by the features:

$$w_x \frac{(x-\mu_x)^2}{2\sigma_x^2}, \; w_y \frac{(y-\mu_y)^2}{2\sigma_y^2}, \; w_l \frac{(l-\mu_l)^2}{2\sigma_l^2}, \; w_c \frac{(c-\mu_c)^2}{2\sigma_c^2} \quad \text{and} \quad w_h \frac{(h-\mu_h)^2}{2\sigma_h^2} \tag{2}$$

where $(l, c, h)$ are the LCH coordinates of each pixel, (x,y) are the relative position of the pixel in the face box, scaled between zero and one, $\mu_i$ and $\sigma_i$ are the mean and standard deviation of each of the $i = l, c, h, x, y$ features in this face box, and $w_l$, $w_c$, $w_h$, $w_x$, $w_y$ are the weights of each of the features, for this specific face, as described in the previous paragraph.

The second learning method assumes that the skin probability model is given by a gaussian distribution in the five dimensional feature space.

$$P(skin|l, c, h, x, y) = Z \cdot e^{-\left(w_l \frac{(l-\mu_l)^2}{2\sigma_l^2} + w_c \frac{(c-\mu_c)^2}{2\sigma_c^2} + w_h \frac{(h-\mu_h)^2}{2\sigma_h^2}\right)} \cdot e^{-\left(w_x \frac{(x-\mu_x)^2}{2\sigma_x^2} + w_y \frac{(y-\mu_y)^2}{2\sigma_y^2}\right)} \tag{3}$$

where $Z$ is a normalization factor. Having found the relative weight of each of the features $w_l$, $w_c$, $w_h$, $w_x$, and $w_y$ for this specific image, we now use the "skin" tagged pixels only, to find the mean and standard deviation of each of the features. For calculating the skin-face map, which attach a skin likelihood estimation to each of the pixels in the face box, we use the probability function which is described by a five dimensional gaussian probability with the calculated means, standard deviations and weights as estimators to the value of the function parameters. This method is marginally less accurate than the nearest neighbor method, but is computationally efficient. In particular, it is fast enough for a production imaging system.

Now that the face-skin regions have been identified, we use it to estimate the specific skin colors for each of the people in the image. Once again, we assume gaussian color model of the color of the skin pixels. Outside the face box the location features are no longer meaningful. The luminance feature is also problematic, as different body regions may have very different luminance conditions, hence we are left with the color features only. The relative importance of those two features, $w_c$ and $w_h$, for each captured person, is already known from the previous steps. We extract the $\mu$ and $\sigma$ parameters of this person's gaussian skin model from the face-skin regions of this person. Each face box yield one skin likelihood function.

$$P(skin|c, h) = Z \cdot e^{-\left(w_c \frac{(c-\mu_c)^2}{2\sigma_c^2} + w_h \frac{(h-\mu_h)^2}{2\sigma_h^2}\right)} \tag{4}$$

The likelihood that a pixel in the image describe skin is the maximal skin-likelihood over all skin-models in the image.

# 3 Applications

## 3.1 Selective Correction of Skin Colors

The adaptive skin map presented in section 2.3 enables automatic selective color correction to image skin tones, where the colors of pixels will be shifted towards the "memory prototype" skin colors, while the magnitude of the shift will be the function of the pixel's likelihood of belonging to the skin tone area.

The color correction is done in the LAB color space, which is the Cartesian representation of the cylindric $LCH$ coordinate,

$$\begin{aligned} a &= c \cdot \cos(h) \\ b &= c \cdot \sin(h) \end{aligned} \tag{5}$$

Let $\bar{a}^*_{ref}$ and $\bar{b}^*_{ref}$ be the skin color memory prototype in the plane of constant lightness in a uniform color space such as CIECAM-UCS [13]. Likewise, $\bar{a}^*_{image}$ and $\bar{b}^*_{image}$ are the mean $a^*$ and $b^*$ coordinates of the image memory color computed using the personal skin map (see Section 2.3). Let $\bar{a}^*_{orig}$ and $\bar{b}^*_{orig}$ be the color coordinates of the pixel. For every pixel in the image, the initial shift $\Delta I$ that is required to move the color towards the prototype one is defined as $\Delta I_a = \bar{a}^*_{ref} - \bar{a}^*_{image}$ and similarly for $b^*$. Then, the new 'corrected' $a^*$ values are computed as:

$$\begin{aligned} a^*_{new} &= a^*_{orig} + \Delta I_a k P \left(a^*_{orig}, b^*_{orig}\right)^\gamma \\ b^*_{new} &= b^*_{orig} + \Delta I_b k P \left(a^*_{orig}, b^*_{orig}\right)^\gamma \end{aligned} \tag{6}$$

Here $P(a^*_{orig}, b^*_{orig})$ is the pre-calculated skin likelihood of this pixel. $\gamma$ controls the smoothness of the transition between the corrected and uncorrected regions, by controlling the extent to which colors with less-then 1 likelihood are shifted. $k$ is a factor that controls the magnitude of color correction, which is required in order to keep the colors "in context" with the image: overcorrected skin tone might perfectly fit the memory prototype, but might look out of place in a particular image's color scheme. We use $\gamma = 2.2$, and $k$ is usually in the range $\frac{1}{4}$ to 1. Example plots of skin tone color correction are shown in Figures 1 and 5.

## 3.2 Skin aware automatic image sharpening

As a rule of thumb people prefer sharp images, so long as they are not over-sharpened. In particular the details of an image should be be bright and sharp. Many home photos are slightly out of focus or blurry, and may often be improved by sharpening. In addition, printing generally induces a blurring function on the image due to half-toning and other reasons, so some sharpening of images before printing can pre-compensate for the later blurring process. However sharpening an image which is already sharp is undesired as it might lead to over-sharpening. But the sharpness of the original image is not the only factor that needs to be taken into account, in a fall proof image enhancement application. Images with hight noise level or compression artifact, e.g. JPEG artifacts, are often better left unsharpened, since sharpening enhance the undesired artifacts. The proposed automatic image sharpening takes those considerations into account and calculates an image specific sharpening parameter (ISSP) base on image sharpness estimation [14], noise estimation [15] and JPEG artifact estimation [16].


(a) (b) (c)

**Figure 5:** *Example of the result of automatic skintone correction. (a) Arrow plot (in CIECAM02-UCS $a^*b^*$ plane) describing the correction applied on the sample image; the beginning of the arrow represents the original image color, arrow head represents the corrected color. (b) Original image. (c) Corrected image. CopyrightNotice : Creative Commons 2.0 / Attribution, http://www.flickr.com/photos/shaggypaul/.*

However, even when the image is good, with no noise or JPEG artifacts, such as Figure 2a, a global image sharpening that does not take the image content into account, shown in Figure 2(b), enhances image features such as the pearl and lace to a pleasing level, but also enhances the wrinkles and pimples on the girl's skin. In skin aware automatic image sharpening we multiply the ISSP by a factor, which differs from pixel to pixel and is a linear function of the skin likelihood of this pixel. Thereby we can sharpen all non-skin regions but keep skin regions unsharpened. In particular, skin regions on face, including wrinkles and skin defects, would not be sharpened while eyes are sharpened to create a happy vivid appearance. Figure 2(c), shows the results of the skin aware automatic image sharpening, when applied on Figure 2(a). Note that the necklace, lace and eyes are vivid and sharp, yet the wrinkles and pimples are not enhanced.

# 4 Evidence that the algorithms work

In this paper, we have presented a novel image-enhancement method that achieves better results by using the face information to identify skin regions. We created a person specific skin map, adaptive to each person in each specific image, and used the skin likelihood information to guide the color correction algorithm. The proposed algorithms were tested at many levels, from the skin map itself to overall image improvement, as described in the following sections.

## 4.1 Skin maps agree with human definitions

Testing the skin maps is challenging, due to the difficulty of tagging all the pixels in the image as skin/non-skin. We overcome this difficulty by pre-segmenting the images, using the segmentation algorithm from [17], and tagging skin segments rather than single pixels. Naturally, the judgment per segment differs than the judgment per pixel. Furthermore segments from an automatic segmentation do not always line up perfectly with skin map. In addition by moving from pixels to segments we compromise the fine differentiation of the skin maps, e.g. segmenting skin regions from eyes, and become dependent in the original segmentation. In that sense the results of this experiment provides only a coarse assessment of the accuracy of our skin maps.

Our test set contained two sets of outdoor images. The first set was used for training and algorithm tuning. This set contained 200 images from the Berkeley segmentation dataset [18], and was tagged by us. The second set was used for testing. This set contained 196 images, and that was tagged by an unbiased observer, i.e., not one of the authors. We found that 80% of the skin pixels were tagged correctly, and 21% of the pixels detected as skin where false detections. The detection results are comparable to the results obtained in a large scale experiment presented in [19], and false detections are somewhat lower. In their experiment, they use a skin color model and non-skin color model and some size and shape considerations and report about 83% detection rate of images with people, with about 30% false detection.

## 4.2 Skin-sensitive color correction yields preferred results

The test for the skin sensitive automatic color correction includes 13 images with varied quality. We printed two copies of each image, one with skin sensitive automatic color correction and the other without it. While some of the images required significant color improvement, the difference in most images was subtle and in two or three images the difference was hardly noticeable. Ten people with no prior knowledge in image processing or of the nature of the tested algorithm, were brought one by one into a room where the pairs of images were presented under controlled lighting conditions. Each person was first asked to point to his favorite image in each of the pairs. Some subjects could not see any difference between the images, and were advised to make initial or random pick. After picking his preferred image from each pair, the person was led to look at the skin colors and, once again, point on his favorite image in each of the pairs.

Figure 6 details test results. Most of the people preferred the color-corrected images. In some images, e.g. image number nine, the preference is small. Those are the images where the difference was hardly noticable. However, after the subjects where led to look for difference in the skin color, most subjects preferred the corrected images.

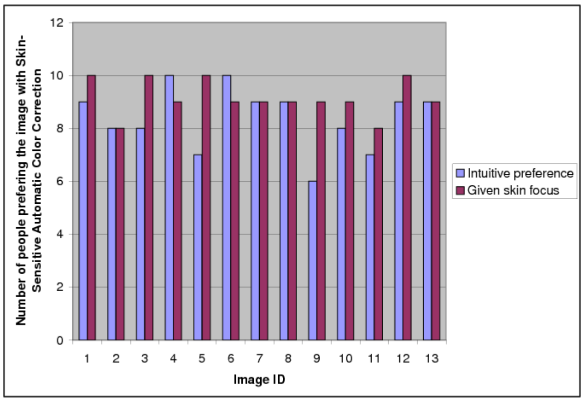


**Figure 6:** *Perceptual test results for the skin sensitive automatic color correction. Ten people were presented with pairs of images and pointed on the image they preferred in each pair. The subjects preferred the images with skin-sensitive color correction over the alternative.*

## 4.3 Qualification tests for the commercial printing system

The algorithms presented in this paper are implemented in a HP-Indigo photo Enhancement Server (HIPIE), which processes home photos prior to printing. The qualification tests for this system included the review of thousands of images, to verify that the color correction improves the image, or, at least, does not harm it [20]. In addition, two Human Metrology tests were performed [21], one compared the system results with the original image, and the other compared the results of system with a previous version, version one.

The previous versions of the system supports skin maps which are a bi-linear combination of the global skin presented in Section 2.2 and a fuzzy face map which is based on the face detection algorithm described in Section 2.1. Version one contains Skin-aware sharpening, but the quality of the skin maps is not sufficient to support skin-aware color correction. Version one is was shown to improve the original image quality by review of hundreds pairs of images, standard side by side methodology tests and by customers feedback.

For images with dominant skin colors, the main differences between the tested system and the previous version are the improved skin maps and the additional skin-selective color correction, hence a sub-test with skin-colored set of images was performed to specifically test those proposed upgrades. The skin colored sub-set of images contains twelve images, most of which contain people and faces. In this test, the average human preference was in favor of the proposed upgrades. In the test where system results were compared to the original images, the mean human preference favored system results for all of the images. When compared with the previous version, the upgraded system was preferred for eight out of the twelve images.

## 5 Conclusions

This paper presented an algorithm for detecting skin in natural images using information from face detection, a global skin color model and machine learning. This algorithm learns a color model for each person in the image, and uses these models to compute a skin map in which every pixel is assigned a probability of being skin.

The skin map is used for low risk skin-sensitive image enhancement. We presented two such applications: color correction and image sharpening. These algorithms are implemented in a real-world imaging production system. The requirements for such a system motivated several main aspects of our approach

- fully automatic processing,
- an overwhelming preference of the enhance image over the original, and
- fast computation

Based on several human metrology tests, we have attained these objectives.

## Acknowledgements

We would like to thank Renato Keshet for initiating, promoting and supporting the object detection research, Shlomo and his team for supporting the HIPIE system, Maya Shalev and Noa Erlich for allowing us to use the data from their report, Sagi Schein and Mani Fisher for their suggestions and remarks, Shimrit Feldman for tagging the testing set, and Mary Nielsen for pointing that state-of-the-art automatic image enhancement sometimes damage faces. We would like to thank Flickr for the use of their images in the paper.

## 6 References

[1] T. Matsui, K. Arakawa, and K. Nomoto, "A nonlinear filter system for beautifying face images with enhancement using interactive evolutionary computing," in *Intelligent Signal Processing and Communications*, 2006, pp. 534 – 537.

[2] K. Kryszczuk and A. Drygajlo, "Color correction for face detection based on human visual perception metaphor," in *Proc. of the Workshop on Multimodal User Authentication*.

[3] Y. Eisenthal, G. Dror, and E. Ruppin, "Facial attractiveness: Beauty and the machine," *Neural Computation*, vol. 18, no. 1, pp. 119–142, 2006.

[4] T. Arici and S. Dikbas, "Skin-aware local contrast enhancement," in *Image Processing, 2007. ICIP 2007. IEEE International Conference on*.

[5] R. Hunt, "The preferred reproduction of blue sky, green grass and caucasian skin in colour photography," *The Journal of Photographic Science*, vol. 22, pp. 144–149, 1974.

[6] P. Bodrogi and T. Tunde, "Colour memory for various sky, skin, and plant colours: effect of the image context," *Color Research and Application*, vol. 26, no. 4, pp. 278–289, 2001.

[7] M. Yamamoto, Y.-H. Lim, X. Wei, M. Inui, and H. Kobayashi, "On the preferred flesh colour in japan, china and south corea," *Imaging Science Journal*, vol. 51, no. 3, pp. 163–174, 2003.

[8] S. R. Fernandez and M. D. Fairchild, “Observer preferences and cultural differences in color reproduction of scenic images,” in *IS&T/SID Tenth Color Imaging Conference*, 2002.

[9] R. Hunt, “How to make pictures and please people,” Colour Image Science: Exploiting Digital Media, pp. 421–437, 2002.

[10] R. Keshet, C. Staelin, B. Oicherman, M. Fischer, P. Kisilev, D. S. Sagi Schein, M. Vans, H. Nachlieli, R. Bergman, R. Maurer, M. Aharon, A. Berkovich, R. Waidman, O. Bengigi, G. Amit, S. Harush, A. Asher, H. Chao, D. Greig, M. Gaubatz, and S. Simske, “Automatic photo enhancement server (hipie 2),” in *IS&Ts Second International Symposium on Technologies for Digital Fulfillment*, 2009.

[11] P. Viola and M. Jones, “Robust real-time face detection,” *International Journal of Computer Vision*, vol. 57, no. 2, pp. 137–154, 2004.

[12] B. O. Martinez, “Semantic retrieval of memory color content,” Ph.D. dissertation, Universitat Autonoma de Barcelona, 2004.

[13] G. C. Luo Ronnier and C. Li, “Uniform colour spaces based on ciecam02 colour appearance model,” *Color Research and Application*, vol. 31, no. 4, pp. 320–330, 2006.

[14] D. Shaked and I. Tastl, “Sharpness measure – towards automatic image enhancement,” in *ICIP’05*, September 2005.

[15] P. Kisilev, D. Shaked, and S. Lim, “Noise and signal activity maps for better imaging algorithms,” in *ICIP*.

[16] R. Maurer and C.-H. Staelin, “Processing an input image to reduce compression-related artifacts,” US patent application, Filed Sept 26, 2007, 2007.

[17] L. G. Ugarriza, E. Saber, V. Amuso, M. Shaw, and R. Bhskar, “Automatic image segmentation by dynamic region growth and multiresolution merging,” *Journal of Image Processing*, vol. 99, no. 99, 2009.

[18] D. Martin, C. Fowlkes, D. Tal, and J. Malik, “A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics,” in *Proc. 8th Int’l Conf. Computer Vision*, vol. 2, July 2001, pp. 416–423.

[19] M. Jones and J. Rehg, “Statistical color models with application to skin detection,” in *In Proceedings IEEE Conference on Computer Vision and Pattern Recognition*, vol. 1, 1999, pp. 274–280.

[20] Shlomo Harush, Nitzan Sapir and Alon Asher, Color and Algorithms Research and Development Team, HP-Indigo.

[21] “Evaluate relative quality of hipie2 to hipie1 and to no hipie,” Maya Shalev and Noa Erlich, HP Indigo Human Metrology lab.