

An Algorithm for Nudity Detection

Rigan Ap-apid
College of Computer Studies
De La Salle University
Manila, Philippines
apapidr@dlsu.edu.ph

ABSTRACT

This paper presents an algorithm for detecting nudity in color images. A skin color distribution model based on the RGB, Normalized RGB, and HSV color spaces is constructed using correlation and linear regression. The skin color model is used to identify and locate skin regions in an image. These regions are analyzed for clues indicating nudity or non-nudity such as their sizes and relative distances from each other. Based on these clues and the percentage of skin in the image, an image is classified nude or non-nude.

The skin color distribution model performs with 96.29% recall and 6.76% false positive rate on a test set consisting of 2,303,824 manually labeled skin pixels and 24,285,952 manually labeled non-skin pixels. The Nudity Detection Algorithm is able to detect nudity with a 94.77% recall and a false positive rate of 5.04% on a set of images consisting of 421 nude images and 635 non-nude images.

Keywords: *pornography, nudity, non-nudity, skin color, correlation, linear regression, skin regions*

1. INTRODUCTION

The rapid development in the field of digital media has exposed us to huge amounts of non-textual information such as audio, video, and images. In fact, one can be easily overwhelmed by the amount of information available through electronic means. More often than not, we are exposed to information that is irrelevant for our own purposes. Thus, retrieval and classification techniques have been and are continuously being developed and improved to facilitate the exchange of relevant information. Correspondingly, there is a clear necessity for user-customized information selection.

In recent years, problems with the lack of control and regulation over what information can be made available, especially over the Internet, have increasingly been apparent. One such problem is the proliferation of pornography or nudity, which has been the focus of various studies in Internet security.

Determining image content continues to be an active research area even as the distribution of all types of images becomes more widespread than ever. Nudity detection is recognized as an important step towards limiting the proliferation of unwanted and damaging information. Technology companies devoted to image recognition acknowledge that the state of the art in nudity detection is still crude, yielding inexact results at the cost of computing power. Thus, new methods for detecting nudity are continually being sought.

This paper describes an algorithm for detecting nudity in images. Contrary to most detection systems, images are not compared to pre-classified images in a database. Aside from raw RGB values, the system uses two other color spaces (HSV and normalized RGB) to aid in the classification. The set of images used is limited to color images in JPEG (Joint Photographic Experts Group) format. However, the detection scheme can theoretically be easily modified to allow for classification of images in other formats.

2. BACKGROUNDS

A basic nudity filtering solution requires the analysis of an image to determine the presence of objects in the image that may signify nudity or inappropriate content. Image features are initially extracted and analyzed, and then classified according to some parameters. Common features that give clues to image content include color, texture, filenames, image dimensions, object shapes, and the number of relevant objects found in the image.

2.1. Skin Color

The most important feature that provides clues to image content is color. Color is a low level feature, which makes it computationally inexpensive and therefore suitable for real-time object characterization, detection and localization (Martinkauppi, 2002).

Nudity often consists of showing naked persons, special shots of sexual organs, or a picture of sexual intercourse (Lin et al., 2003). These images show a lot of skin and thus, skin color is a basic feature used in nudity detection. A disadvantage of systems using color as a primary feature is that the systems will not work with black and white images. However, nude images are rarely in black and white.

The simplest methods in skin detection define or assume skin color to have a certain range or values in some coordinates of a color space. Usually, the skin is recognized due to the fact that it is composed of a group of color containing red for the blood and yellow and brown for melanin.

In general, the fairness or darkness of skin depends on the amount of melanin in the skin. Due to the emerging importance of skin detection in computer vision several studies have been made on the behavior of skin chromaticity at different color spaces. Many studies such as those by Yang and Waibel (1996) and Graf et al. (1996) indicate that skin tones differ mainly in their intensity value while they are very similar in chrominance coordinates. Several surveys on the use of color spaces such as those by Zarit et al. (1999) and Vezhnevets et al. (2003) have been made to identify optimal methods for skin detection using different color spaces. While some studies show differences in performance due to the color space used (Terillon et al., 2000), (Vezhnevets et al., 2003), some researchers (Bosson et al., 2002), (Chan et al. 1999) agree that the choice of color space is not critical provided that enough training is done and that an appropriate data set is used.

The main goal of skin color detection or classification is to build a decision rule that will discriminate between skin and non-skin pixels. Identifying skin colored pixels involves finding the range of values for which most skin pixels would fall in a given color space. This may be as simple as explicitly classifying a pixel as a skin pixel if $R > G$ or $R > B$ or both (Brown et al., 2000) or may be as complex as models using neural networks and Bayesian methods (Jones & Rehg, 1999), (Chai & Bouzerdoum, 1999). In general, a good skin color model must have a high detection rate and a low false-positive rate. That is, it must detect most skin pixels while minimizing the amount of non-skin pixels classified as skin.

2.2. Color Spaces

The purpose of a color space is to facilitate the specification of colors in some standard, generally accepted manner. A color space is a specification of a coordinate system and subspace within a system where each color is represented by a single point. Most color spaces today are oriented toward hardware such as color monitors or toward applications where color manipulation is a goal such as the creation of color graphics or animation (Gonzalez & Woods, 2002).

Various color spaces are used for processing digital images. For some purposes, one color space may be more appropriate than others. For skin detection, researchers do not quite agree on whether the choice of color space is critical to the overall performance of the detection system. One obvious

reason for this is the lack of a standard set of images that can be used as a benchmark for algorithms using different color spaces.

2.2.1. The RGB Color Space

The RGB color space originated from CRT display applications. In the RGB space each color appears in its primary spectral component of red, green, and blue. Images represented in the RGB space consist of three component images, one for each primary color. When fed into an RGB monitor, these images combine on the phosphor screen to produce a composite color image (Gonzalez & Woods, 2002).

The RGB color space is one of the most widely used color spaces for storing and processing digital image. However, the RGB color space alone is not reliable for identifying skin-colored pixels since it represents not only color but also luminance. Skin luminance may vary within and across persons due to ambient lighting so it is not dependable for segmenting skin and non-skin regions. Chromatic colors are more reliable and these are obtained by eliminating luminance through some form of transformation. The color spaces Normalized RGB, HSV, and YCbCr are transformations commonly used by studies on skin color (Waibel et al., 1999).

2.2.2. The HSV (HSI, HSL) Color Space

The RGB color space is ideal for hardware implementations. However, it is not well suited for describing colors in terms that are practical for human interpretation. For example, we do not refer to the color of a car in terms of the percentages of each of the primary components composing color. We also do not think of a color image as being the composite of three images. When humans view color objects, we describe them in terms of hue, saturation, and brightness (Gonzalez & Woods, 2002).

The HSV (Hue, Saturation, Value/Intensity/Luminance) color space describes color with intuitive values, based on the artist's idea of tint, saturation and tone. This was introduced when there was a need to specify color properties numerically. *Hue* defines the dominant color as described by wavelength, for instance the distinction between red and yellow. *Saturation* measures the colorfulness of an area in proportion to its brightness such as the distinction between red and pink. *Value* refers to the color luminance, the distinction between a dark red and a light red.

For skin detection, the value component is discarded to eliminate the undesirable effect of uneven illumination. The transformation is defined by

$$H = \arccos \frac{\frac{1}{2}((R-G) + (R-B))}{\sqrt{((R-G)^2 + (R-B)(G-B))}}$$

$$S = 1 - 3 \frac{\min(R, G, B)}{R + G + B}$$

$$V = \frac{1}{3}(R + G + B)$$

Some studies show that HSV is invariant to highlights at white light sources, to matte surfaces, and ambient lighting. However, hue discontinuities and the computation of the luminance component conflict badly with the properties of color vision. The cyclic nature of Hue-Saturation spaces also makes it inconvenient for parametric skin color models that need a tight cluster of skin performance for optimum performance (Vezhnevets et al., 2003).

2.2.3. The Normalized RGB Color Space

The Normalized RGB color space removes the luminance component through the normalization

$$r = \frac{R}{R + G + B}$$

$$g = \frac{G}{R + G + B}$$

$$b = \frac{B}{R + G + B}$$

Since the sum of the three components r, g, and b is 1, the blue (b) component can be omitted to reduce space dimensionality.

2.3. Classifying Images

After extracting the relevant features of an image, filtering systems use various algorithms to classify the image as pornographic (nude) or otherwise. Since pornography is a relative term, most filtering systems have sensitivity settings to allow the user to have some degree of control over what images are classified as pornographic or not. Nevertheless, some images may be wrongly classified. Nude images classified as non-nude are referred to as false-negatives while non-nude images classified as nude are called false-positives. The efficiency of filtering systems can be measured in terms of the amount of false-negatives or false-positives.

In classifying images, most systems use skin color models to process and characterize the images. In these systems, the color model is trained on a number of examples under varying degrees of illumination conditions (Jones & Rehg, 1998). A common procedure for image characterization is by comparing an image to a database of pre-classified images using histogram color models.

The classification algorithm by Wang et al. (1997) involves matching an image against a small number of feature vectors obtained from a training database of 500 objectionable images and 8,000 benign images. If the image closely matches two or more of the marked objectionable images in the training database within the closest 15 matches, it is

considered objectionable. The algorithm works with 9% false-positives.

The pornography detection algorithm developed by Lin et al. (2003) uses the correlation between skin and non-skin regions as the input of a support vector machine (SVM), a tool for classification with learning abilities. The algorithm has an accuracy of 75% and a false alarm rate of 35%.

Jones & Rehg (1998) constructed skin and non-skin models from a dataset of nearly one billion labeled pixels. A neural network was then trained on 5453 adult images and 5226 non-adult images. The procedure resulted in 93% correct detections with 8% false-positives. Lin, Chen, & Fuh (2003) contends that the training process for neural networks is very slow and difficult because pornography classification is a matter of opinion.

Pornsweeper™ (Bosson, 2002) first compares an image with images in its *exception list*, which are common images pre-classified by an administrator as pornographic or not. If an image is found to be in the exception list, it is quarantined for administrator inspection. If the image is not on the list, it is passed on to an analyzer for further analysis. Pornsweeper™ detects nude images with 15% false negatives.

We can compare the performance of nudity detection systems in terms of the percentages of the images correctly or wrongly classified by the systems. However, this method is very unreliable because different systems use different test data sets (Jones & Rehg, 1998). Different data sets generally yield different results.

3. THE ALGORITHM

The Nudity Detection Algorithm is based primarily on observations that in general, nude images contain large amounts of skin, people have different skin tones, and skin regions in nude images are relatively close to each other. These observations require the identification of skin and non-skin pixels in an image to be classified. The identified skin pixels are analyzed to determine which among them are connected or form continuous regions. These skin regions are further analyzed for clues of nudity or non-nudity.

In general, the Nudity Detection Algorithm consists of the following steps:

1. Detect skin-colored pixels in the image.
2. Locate or form skin regions based on the detected skin pixels.
3. Analyze the skin regions for clues of nudity or non-nudity.
4. Classify the image as nude or not.

The first step in the algorithm requires a skin color distribution model, which is the basis for classifying pixels as skin or non-skin.

Most skin color distribution models use a single color space to classify pixels as skin-colored or not (Chai & Bouzerdoum, 1999), (Lin et al., 2003), (Forsyth & Fleck, 1999), (Jones & Rehg, 1999). However, based on the researcher's experiments, the performance of a model can be improved by using two or more other color spaces. Thus, aside from RGB values, the skin color model used in this system uses the Normalized RGB and HSV color spaces.

The set of images used for developing the color model in this study consists of manually labeled skin regions representing all types of skin color. Contrary to some studies, different skin tones do not exhibit the same chromaticity if the luminance component is removed. Thus, the set of images were partitioned into three subsets based on perceived illumination level – light, brown, and dark. Using correlation and linear regression the color properties of each subset are analyzed and used to develop the skin color model.

Experiments show that the skin filter developed for this study works with 96.29% recall and a 6.76% false positive rate on the training set.

The percentage of skin pixels relative to the image size is one of the main features that affect the probability of an image being nude or not. The larger the percentage of skin is, the higher the probability that an image will be classified nude. Experiments show that a threshold of 15% yields good results in terms of recall. Thus, for this study an image is considered not nude if it contains less than 15 percent skin. If the skin percentage is greater than or equal to 15, the image has to be analyzed further.

The next step in the algorithm is the analysis of detected skin pixels in terms of adjacency or continuity to identify or locate skin regions. The number of skin regions, the sizes of the three largest skin regions and the relative positions of these skin regions are important features that aid in the classification.

Although the percentage of skin is a good indicator of the presence or absence of nudity in an image, it allows too many false positives if used alone. The results show that the performance can be improved by combining skin percentage feature and region properties.

In summary, the nudity detection algorithm works in the following manner:

1. Scan the image starting from the upper left corner to the lower right corner.
2. For each pixel, obtain the RGB component values.
3. Calculate the corresponding Normalized RGB and HSV values from the RGB values.
4. Determine if the pixel color satisfies the parameters for being skin established by the skin color distribution model.
5. Label each pixel as skin or non-skin.
6. Calculate the percentage of skin pixels relative to the size of the image.
7. Identify connected skin pixels to form skin regions.
8. Count the number of skin regions.

9. Identify pixels belonging to the three largest skin regions.
10. Calculate the percentage of the largest skin region relative to the image size.
11. Identify the leftmost, the uppermost, the rightmost, and the lowermost skin pixels of the three largest skin regions. Use these points as the corner points of a bounding polygon.
12. Calculate the area of the bounding polygon.
13. Count the number of skin pixels within the bounding polygon.
14. Calculate the percentage of the skin pixels within the bounding polygon relative to the area of the polygon.
15. Calculate the average intensity of the pixels inside the bounding polygon.
16. Classify an image as follows:
 - a. If the percentage of skin pixels relative to the image size is less than 15 percent, the image is not nude. Otherwise, go to the next step.
 - b. If the number of skin pixels in the largest skin region is less than 35% of the total skin count, the number of skin pixels in the second largest region is less than 30% of the total skin count and the number of skin pixels in the third largest region is less than 30% of the total skin count, the image is not nude.
 - c. If the number of skin pixels in the largest skin region is less than 45% of the total skin count, the image is not nude.
 - d. If the total skin count is less than 30% of the total number of pixels in the image and the number of skin pixels within the bounding polygon is less than 55 percent of the size of the polygon, the image is not nude.
 - e. If the number of skin regions is more than 60 and the average intensity within the polygon is less than 0.25, the image is not nude.
 - f. Otherwise, the image is nude.

The threshold values used in the algorithm have been determined empirically.

4. RESULTS

Separate sets of images were collected for training and testing the algorithms developed in this study. The training set for the skin filter consisted of 1,182,608 manually labeled skin pixels and 10,471,553 manually labeled non-skin pixels while the testing set consisted of 2,303,824 manually labeled skin pixels and 24,285,952 manually labeled non-skin pixels. The training set for the nudity detection algorithm consisted of 264 nude images and 300 non-nude images. The testing set consisted of 421 nude images and 635 non-nude images.

4.1. The Performance of the Skin Filter

Experimental results suggest that there is no significant difference between the performance of the skin filter on the training and testing images. The skin filter performs with a 97% correct detection rate and an 8% false positive rate on the testing set.

The skin color distribution model was compared to several common models using the set of images obtained for this study. The results are shown in Table 1. As shown in the table the skin filter developed and used for this algorithm outperforms the other skin filters.

Table 1. A Comparison of Skin Color Models

Skin Filter	Correct Detection (%)	False Positives (%)
Kovac et al., 2003	10	42.5
Marius et al., Lin et al., 2003	34	7.5
Solina et al., 2002	98	14
The Skin Color Distribution Model in This Study	96.29 98	6.76 12.46

4.2. The Performance of the Nudity Detection Algorithm

The performance of the Nudity Detection Algorithm on the testing images is unexpectedly better than its performance on the training images. On the training set, the algorithm worked with a 92.80% recall and an 8.67% false positive rate. On the testing set, the algorithm correctly identified 399 out of 421 nude images, a 94.77% recall. It wrongly classified 32 out of 635 non-nude images, a false positive rate of 5.04%.

False alarms cannot be avoided because some benign images have characteristics that are very similar to that of nude images. Some images are wrongly classified as nude due to the presence of features usually associated to nudity. Possible reasons are the presence of large objects whose colors are very similar to skin and clothed or scantily clad persons resulting in continuous skin regions being detected.

The researcher was not able to obtain data sets used by other researchers dealing with nudity detection. The researcher was also unable to obtain detailed algorithms or software implementations of the algorithms. Thus, it is impossible to compare the algorithms directly using the same image set. The reported results can be compared but a meaningful evaluation of the relative performances cannot be done. Table 2 shows a comparison of the Nudity Detection System with other similar systems in terms of the reported results.

Table 2. A Comparison With Other Systems

Algorithm/System	Recall (%)	False-Positives (%)
SVM	86	35
WIPE	97.5	18.4
Jones-Rehg	86.7	9
Fleck-Forsyth	43	4.2
The Nudity Detection Algorithm	94.32	5.98

Although it cannot be concluded that the Nudity Detection Algorithm designed by the proponent is better than other similar systems, it can be said that it performs relatively well compared to other similar algorithms.

5. CONCLUSIONS

The proponent is inclined to agree with the belief that the choice of color space is not crucial as shown by several studies. However, using a single color space may limit the performance of the skin color filter. This study shows that better performance can be achieved by using two or more color spaces.

Contrary to some studies, experimental results in this study suggest that chrominance values still significantly differ for light and dark skin even if the luminance component is removed through a normalization or transformation of RGB values. Thus, a skin color model should not treat light and dark skin in the same manner. This may be more computationally expensive but it allows for more skin pixels to be labeled correctly.

The proponent was unable to find researches on skin color using linear regression. The correlation coefficients for the color values (Normalized RGB and HSV) for different skin color suggest a high linear correlation. Thus, it is possible to model skin color using linear regression analysis and this is confirmed by the results of this study. However, linear regression may not be the best model for skin color. The relationship between the values may not be strictly linear. If this is the case, using curves may be more appropriate than using lines to model the relationship.

The proponent believes that the performance of the system can be improved by using similarity measures or template matching where images are compared to pre-classified nude images. However, large memory requirements for the use of templates may present undesirable consequences. The appropriateness of the templates used must also be carefully established.

The performance may also be improved by using convex bounding polygons and using polygons with more than four sides. The drawback is the effect on execution time. It is computationally more expensive to identify interior points for concave polygons than for convex ones. Increasing the number of sides also has a corresponding increase in execution time. Thus, care must be taken in determining the appropriate threshold for the optimal number of sides of the bounding polygon.

The evaluation of performance results would have been more meaningful if the comparison with other similar systems was done using a single image set. However, it is difficult to obtain image sets used in other researches on skin color. The implementations of the algorithms are also not publicly available. Thus, it would be worthwhile for someone to develop a benchmark for algorithms in skin and nudity detection. This would enable researchers on the field to establish standard performance targets and to carry out objective performance evaluations.

6. REFERENCES

- [1] Bosson, A., Cawleyz, G., Chanz, Y., & Harveyz, R. (2002). Non-retrieval: blocking pornographic images. [online]. Available: <http://www2.cmp.uea.ac.uk/~rwh/research/reprints/civr2002.pdf>. (17 September 2004)
- [2] Brown, D., Craw, I., & Lewthwaite, J. A SOM Based Approach to Skin Detection with Application in Real Time Systems. In *Proc. Of the British Machine Vision Conference*, 2001.
- [3] Chan, Y, Harvey, R., & Smith, D. Building Systems to Block Pornography. In *Proc. Of Challenge of Image Retrieval, BCS Electronic Workshops in Computing Series*, pp.34-40, 1999.
- [4] Chai, D. & Bouzerdoum, A. A Bayesian Approach to Skin Color Classification in YCbCr Color Space. In *Proc. Of IEEE Region Ten Conference*, vol. 2, 421-4124.
- [5] Forsyth, D.A. & Fleck, M.M. (1999). Automatic Detection of Human Nudes. In *Proc. International Journal of Computer Vision*, 32(1), 63-77.
- [6] Gomez, G. On Selecting Colour Components for Skin Detection. In *Proc. Of the ICPR*, vol. 2, 961-964, 2000.
- [7] Gomez, G & Morales, E. Automatic Feature Construction and a Simple Rule Induction Algorithm for Skin Detection. In *Proc. Of the ICML Workshop on Machine Learning in Computer Vision*, 31-38, 2002.
- [8] Jones, M. & Rehg, J. (1998) Statistical Color Models with Application to Skin Detection. [online] Available: <http://www.hpl.hp.com/techreports/Compaq-DEC/CRL-98-11.pdf>. (30 September 2004)
- [9] Kovac, J. (2002). Eliminating the Influence of Nonstandard Illumination From Images. Diploma Thesis, University of Ljubljana.
- [10] Lin, Y., Tseng, H. & Fuh, C. Pornography Detection Using Support Vector Machine. 16th IPPR Conference on Computer Vision, Graphics and Image Processing (CVGIP 2003). Kinmen, ROC.
- [11] Martinkauppi, B. 2002. Face Colour Under varying Illumination. [online]. Available: <http://herkules.oulu.fi/isbn9514267885/isbn9514267885.pdf>. (02 October 2003)
- [12] Peer, P., Kovac, J, & Solina, F. Human Skin Colour Clustering for Face Detection. In *Proc. Of EUROCON 2003 – International Conference on Computer as a Tool*, 2003.
- [13] Pornsweeper Review (2002). [online]. Available: <http://www.dansdata.com/pornsweeper.htm>. (24 September 2003)
- [14] Solina, F., Peer, P., Batagelj, B., & Juvan, S. 15 seconds of Fame – An Interactive, Computer-Vision Based Art Illustration. In *Proc. Of the 7th Int'l Conference on Control, Automation, Robotics, and Vision*, pp. 198-204, 2002.
- [15] Terrillon, J., Shizarit M, Fukamachi, H, & Akamatsu, S. Comparative Performance of Different Skin Chrominance Models and Chrominance Spaces for the Automatic Detection of Human Faces in color Images. In *Proc. Of the International Conference on Face and Gesture Recognition*, pp. 54-61, France, 2000.
- [16] Vezhnevets, V., Sazonov, V. & Andreeva, A. A Survey on Pixel-Based Skin Color Detection Techniques. In *Proc. Graphicon-2003*, pp. 85-92, Moscow, Russia, Sep. 2003.
- [17] Waibel, A., Yang, J., & Lu, W. Skin-color modeling and Adaptation. In *Proc. Of ACCV1998*, 687-694, 1998.
- [18] Wang, J., Li, J., Wiederhold, G., & Firschein, O. System for Screening Objectionable Images Using Daubechies' Wavelets and Color Histograms. In *Proc. of the International Workshop on Interactive Distributed Multimedia Systems and Telecommunication Services*, pp. 10-30, 1997.
- [19] Zarit, B., Super, B. & Quek, F. Comparison of Five Color Models in Skin Classification. In *Proc. Of ICCV '99 International Workshop on Recognition, Analysis, and Tracking of Faces and Gesture*