

# A Self-adaptation Method for Human Skin Segmentation based on Seed Growing

Anderson Carlos Sousa e Santos and Helio Pedrini

*Institute of Computing, University of Campinas, Campinas-SP, 13083-852, Brazil*

**Keywords:** Skin Detection, Image Analysis, Face Detection, Color Models.

**Abstract:** Human skin segmentation has several applications in image and video processing fields, whose main purpose is to distinguish image portions between skin and non-skin regions. Despite the large number of methods available in the literature, accurate skin segmentation is still a challenging task. Many methods rely on color information, which does not completely discriminate the image regions due to variations in lighting conditions and ambiguity between skin and background color. Therefore, there is still need to adapt the segmentation to particular conditions of the images. In contrast to the methods that rely on faces, hands or any other body content detector, we describe a self-contained method for adaptive skin segmentation that makes use of spatial analysis to produce regions from which the overall skin can be estimated. A comparison with state-of-the-art methods using a well known challenging data set shows that our method provides significant improvement on the skin segmentation.

## 1 INTRODUCTION

Human skin detection in digital images is a crucial part in several applications, such as face detection, gesture analysis, content-based image retrieval, nudity detection and consequent adult content filtering.

Skin detection can be seen as a classification problem, whose purpose is to determine which image pixels belong to the skin or non-skin classes. Such task presents many challenges. For instance, variation in scene illumination interferes with the appearance of the skin, different cameras produce distinct colors and ethnic diversity promotes various skin tones. These factors make the skin detection process quite complex.

Most studies (Kawulok et al., 2014; Kakumanu et al., 2007; Phung et al., 2005) use color information as evidence to detect skin since this property is able to provide computational efficiency while it demonstrates to be robust to occlusions and rotation and scaling transformations (Kakumanu et al., 2007).

The main obstacle is the existence of an overlap between skin and non-skin colors that occurs independently of the color space. To minimize that, some segmentation methods (Kawulok et al., 2014) attempt to perform a color separability dependent on the skin color that appears in the image. The most significant improvements are performed with the aid of face detection, which provides an estimation of the remaining skin.

This work presents an adaptive segmentation method with no need for face or any other body content detection. It is based on an estimation from regions found through spatial analysis performed with a skin probability map. Experiments conducted on a large well known data set show that our method outperforms other skin segmentation approaches available in the literature.

The text is organized as follows. Section 2 presents the main concepts and work related to skin detection. Section 3 describes the proposed method for self-adaptive segmentation of human skin. Experimental results are shown in Section 4. Finally, Section 5 concludes the paper with final remarks and directions for future work.

## 2 BACKGROUND

Many approaches have been proposed to address detection and segmentation of human skin. The simplest strategy is based on static decision rules defined in some color space or even in a combination of them. Various different color spaces have been explored, such as RGB (Fleck et al., 1996), normalized RG (Soriano et al., 2000), HSV (Sobotka and Pitas, 1998; Wang and Yuan, 2001), YCbCr (Hsu et al., 2002). A transformation of RGB into a single-channel has been proposed by Cheddad et al. (Cheddad et al., 2009), specifically for the purpose of skin detection.

A more sophisticated strategy relies on modeling the statistical distribution of color. There are two main approaches to that: parametric and non-parametric techniques. Both approaches calculate the probability of a given color ( $c$ ) to be skin ( $P(\text{skin}|c)$ ), which generates a probability map such that segmentation can be performed through a threshold. However, parametric approaches assume that the skin distribution fits some explicit model.

Most of the methods available in the literature rely on mixture of Gaussians (Yang and Ahuja, 1999), although there are also methods based on single Gaussian (Subban and Mishra, 2014) or elliptical boundary model (Lee and Yoo, 2002).

On the other hand, non-parametric models estimate the probabilities directly from the training data without any assumptions on its distribution shape. To do so, histograms for skin ( $H_{\text{skin}}(c)$ ) and non-skin ( $H_{\neg\text{skin}}(c)$ ) are built over an annotated data set. The conditional probabilities are obtained as

$$P(c|\text{skin}) = \frac{H_{\text{skin}}(c)}{\sum(H_{\text{skin}}(i))} \quad (1)$$

$$P(c|\neg\text{skin}) = \frac{H_{\neg\text{skin}}(c)}{\sum(H_{\neg\text{skin}}(i))} \quad (2)$$

From these probabilities, Bayes rule can be applied to produce the desired posterior probability as

$$P(\text{skin}|c) = \frac{P(c|\text{skin})P(\text{skin})}{P(c|\text{skin})P(\text{skin}) + P(c|\neg\text{skin})P(\neg\text{skin})} \quad (3)$$

where  $P(\text{skin})$  and  $P(\neg\text{skin})$  are the prior probabilities and usually are set to 0.5. Vezhnevets et al. (Vezhnevets et al., 2003) demonstrated that the choice of the prior probabilities does not influence the overall result. Jones and Rehg (Jones and Rehg, 2002) conducted an extensive evaluation of the method and showed that it outperforms Gaussian mixture models for a sufficient large data set.

There is usually an overlap between skin and non-skin colors when the models become more accurate (Kawulok et al., 2014). Taking this fact into consideration, many researchers have adapted the mentioned methods according to the context. For instance, Kovac et al. (Kovac et al., 2003) defined different rules depending on lighting conditions, whereas Phung et al. (Phung et al., 2003) created an iterative method for determining an optimal threshold for the probability map of a particular image.

Nevertheless, the most significant results are obtained by content-based adaptation, more specifically for face detection. The first of such approaches (Fritsch et al., 2002) uses the region acquired by a face detector to update a unimodal

Gaussian previously determined. Taylor and Morris (Taylor and Morris, 2014) recently proposed to only use the skin of the face in normalized RG to build a Gaussian model, discarding any previous training. A more robust technique (Kawulok, 2010) uses the face region to build a local skin histogram and, consequently, a  $P_{\text{face}}(\text{skin}|c)$  is derived and combined with the general probability for the final map.

Another strategy, namely spatial analysis, considers the structural alignment in the neighborhood of pixels classified as skin - generally with a probability map - such that it refines the segmentation process by removing false positives.

Most of these methods perform an expansion of seeds found by a high threshold. The expansion can be performed through different criteria: threshold hysteresis (Ruiz-del Solar and Verschae, 2004), energy accumulation (Kawulok, 2010) or cost propagation (Kawulok, 2013). The latter one is the more complex and provides superior results, where the Dijkstra's algorithm (Dijkstra, 1959) is used to calculate shortest routes in a combined domain composed of luminance, hue and skin probability.

Some of the methods described in this section will be used for comparison in our experiments, described in Section 4. For a more extensive review of the state-of-the-art methods, refer to Kawulok et al. (Kawulok et al., 2014).

### 3 PROPOSED METHODOLOGY

We propose a method for skin segmentation that combines spatial analysis and adaptive models for better skin probability estimation. The methodology can be divided into three main steps. First, seeds are extracted from the probability map through a precise and systematic strategy for spreading them over the images. Second, a controlled propagation method is applied to grow the seeds into skin blobs. Finally, these blobs are used to estimate the skin color present in the images and such that can optimize the probability map. The main steps of our method are presented in the diagram shown in Figure 1.

Figure 2 illustrates the application of the proposed method to an input image, where the general probability map, the extracted seeds, the blobs after propagation, the final probability map and the resulting segmentation are shown.

Algorithm 1 describes in details the main steps performed by the proposed skin segmentation method.

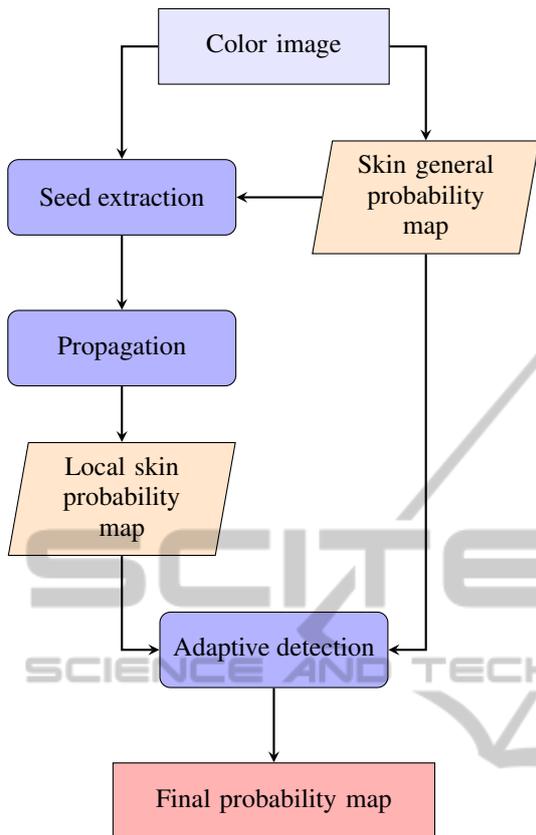


Figure 1: Main stages of the skin detection process.

### 3.1 Seed Extraction

The most important step in a region growing algorithm is the proper choice of the seeds. Since seeds can correspond to false positives, the propagation process does not guarantee that inadequate seeds do not occur in the images, which can compromise the final estimation.

As described in Section 2, skin region propagation methods usually rely on a fixed high threshold and size based analysis for producing the seeds. Such methods do not take different characteristics of each image into account, as well as its respective probability map, once the same threshold is applied to all images. Another disadvantage is the assumption that the resulting skin-like seeds (false positives) are very small, which is not always true. Taking these factors into consideration, we propose an adaptive seed extraction with a homogeneity-based analysis.

In order to obtain the best high threshold for a particular image, we first apply a median filter (Gonzalez et al., 2009) to the probability map and then take the maximum probability (Line 3, Algorithm 1). To allow for images with no skin, if the maximum value is

---

**Algorithm 1:** Proposed skin segmentation method based on seed growing.

---

**input** : color image  $I$ , histogram of skin ( $H_{skin}$ ) and non-skin ( $H_{nonskin}$ ) colors.

**output:** Final probability map  $M_{final}$

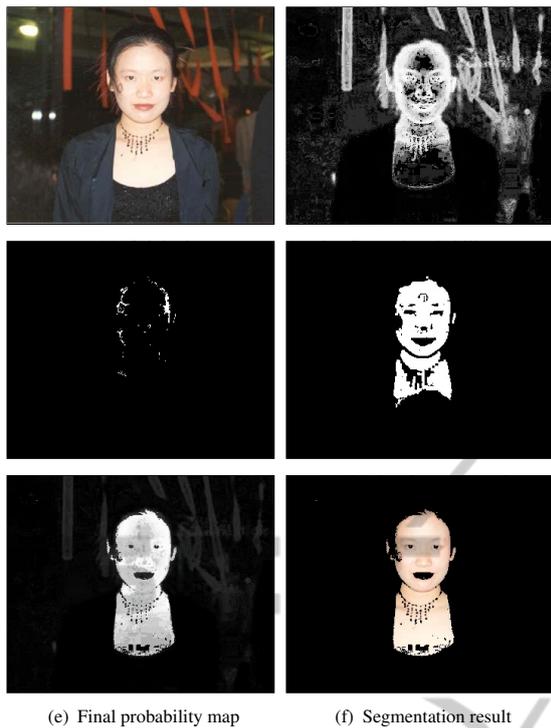
```

1 Build general probability map ( $M_{global}$ )
  according to Equations (1), (2) (3), using  $H_{skin}$ 
  and  $H_{nonskin}$ 
2  $M_{blur} \leftarrow \text{blur}(M_{global}, size)$ 
3  $T_{seed} \leftarrow \max(M_{blur})$ 
4 if  $T_{seed} \leq 0.5$  then
5   return  $M$ 
6 end
7  $edges \leftarrow \text{edgeDetector}(I)$ 
8 for  $x \in I$  do
9   if  $M_{global}(x) \geq T_{seed} \wedge x \notin edges$  then
10     $Seeds \leftarrow x$ 
11   end
12 end
13  $Q \leftarrow \{Seeds\}$  {where  $Q$  is a priority queue}
14 for  $x \in I$  do
15   if  $x \in Seeds$  then
16     $C(x) = 0$ 
17   else
18     $C(x) = -1$ 
19   end
20 end
21 while  $Q \neq \emptyset$  do
22    $q = \text{pop}(Q)$ 
23   for  $s \in \text{Neighbors}(q)$  do {8 -
  neighborhood}
24      $c = C(q) + \rho(q \rightarrow s)$ 
25     if  $(c < C(s) \vee C(s) < 0) \wedge s \notin edges$ 
  then
26        $C(s) = c$ 
27        $Q \leftarrow s$ 
28     end
29   end
30 end
31 Normalize  $C$  by scaling the costs from 0 for the
  maximal cost to 1 for a zero cost
32 for  $x \in I$  do
33   if  $C(x) > 0$  then
34      $H_{skin\_local}(color(x))++$ 
35   end
36 end
37 Build final probability map ( $M_{final}$ ) according
  to Equation (5) using  $H_{skin\_local}$ 
38 return  $M_{final}$ 

```

---

smaller than a minimum threshold ( $T_{seedmin}$ ), it is discarded, otherwise it is assigned as the seed threshold for the original probability map. Therefore, we obtain



(e) Final probability map

(f) Segmentation result

Figure 2: Examples of images obtained by applying each stage of our method to an input image.

seeds with high probability by considering the context.

To prevent the occurrence of false positives, we exclude the choice of seeds located in edge regions, since skin is usually a smooth and homogeneous region. The used edge detector is described in the following section.

### 3.2 Propagation

The objective here is to expand the seeds into skin blobs. The main drawback of propagation methods is the “leakages”, that is, the seed growth to a region of non skin. In order to avoid those, we establish a strict control of the propagation. We modified the cost propagation proposed in Kawulok (Kawulok, 2013) by adding a constraint in which the propagation cannot flow out the image edges. As a consequence, an increase in the false negative rate is concerned with a reduction of the false positive rate. The next step will address these undetected skin regions.

Prior to the actual propagation, an edge detector is applied to the images. In order to benefit from color information, we utilize a color edge detection technique that combines (through logical *or* operator) the results of Canny detector (Gonzalez et al., 2009) for each of the three channels in HSV color space. Fol-

lowing that, a morphological dilation operation is performed, such that small gaps can be closed.

The propagation process is similar to the work by Kawulok (Kawulok, 2013) except that the process stops when an edge is reached for a certain direction (Line 25, Algorithm 1). Besides preventing false positives, this also speeds up the algorithm, once the original approach calculates the costs from the seeds to every other pixel in the image.

### 3.3 Adaptive Detection

Once we have generated the skin blobs, we use them to build a local statistical model (Line 34, Algorithm 1) that adapts to the particular conditions of the image. From the histogram of these resulting skin regions, we obtain a  $P_{local}(c|skin)$ . As for non-skin, we assume that the local distribution follows the global one, which gives

$$P_{local}(c|\neg skin) = P_{global}(c|\neg skin) \quad (4)$$

The final probability is defined as

$$P(skin|c) = \gamma P_{local}(skin|c) + (1 - \gamma) P_{global}(skin|c) \quad (5)$$

where  $P_{local}(skin|c)$  and  $P_{global}(skin|c)$  are both calculated as in Equation (3), differentiating by using local and global data, respectively. The parameter  $\gamma$  controls the importance of the local model.

From Equation (5), we generate the final skin probability map, in which the detection can be performed through a fixed threshold or generated by more complex techniques developed for the general probability map. However, the description of such methods is beyond the scope of this paper.

## 4 EXPERIMENTS

The experiments were evaluated through two different data sets. To train the Bayes classifier, we used 8963 non-skin images and 4666 skin images from the Compaq database (Jones and Rehg, 2002), which contains images acquired from the Internet in a diverse variety of settings and its approximately 1 billion pixels makes it sufficient large for non-parametric estimation of skin color distribution.

For evaluation and comparison purposes, we used the ECU database (Phung et al., 2005) divided into a 1000 images for validation and 3000 images for test. This data set ensures a diversity in terms of background scenes, lighting conditions and skin types.

Both data sets provide a ground-truth that makes possible identify the pixel class (skin or non-skin) for the training and quantitatively evaluate the detection output.

The performance of the skin detection was measured through a number of metrics: true positive rate ( $\eta_{tp}$  - percentage of skin correctly classified as skin); false positive rate ( $\delta_{fp}$  - percentage of non-skin classified as skin); precision ( $\eta_{prec}$  - percentage of correctly classified pixels out of all the pixels classified as skin);  $F_{score}$  (harmonic mean between  $\eta_{prec}$  and  $\eta_{tp}$ ) and detection error ( $\delta_{min} = (1 - \eta_{tp}) + \delta_{fp}$ ). Additionally for non-binary classification, the ROC (receiver operating characteristics) and the respective area under curve (AUC) are applied.

All the following experiments were conducted on an Intel Core i7 3.50GHz with 32GB RAM running 64 bits Ubuntu 12.04 operating system.

In order to determine the bin size of the histogram, we experimented a number of different sizes, as shown in Table 1. Since 32 bins per channel produced the highest value of AUC, this value was used in the proposed method both for local and general models.

Table 1: Bin size evaluation for validation data set.

Bin size	AUC
$8^3$	0.892348
$16^3$	0.918403
$32^3$	<b>0.934036</b>
$64^3$	0.923436
$128^3$	0.917958
$256^3$	0.914159

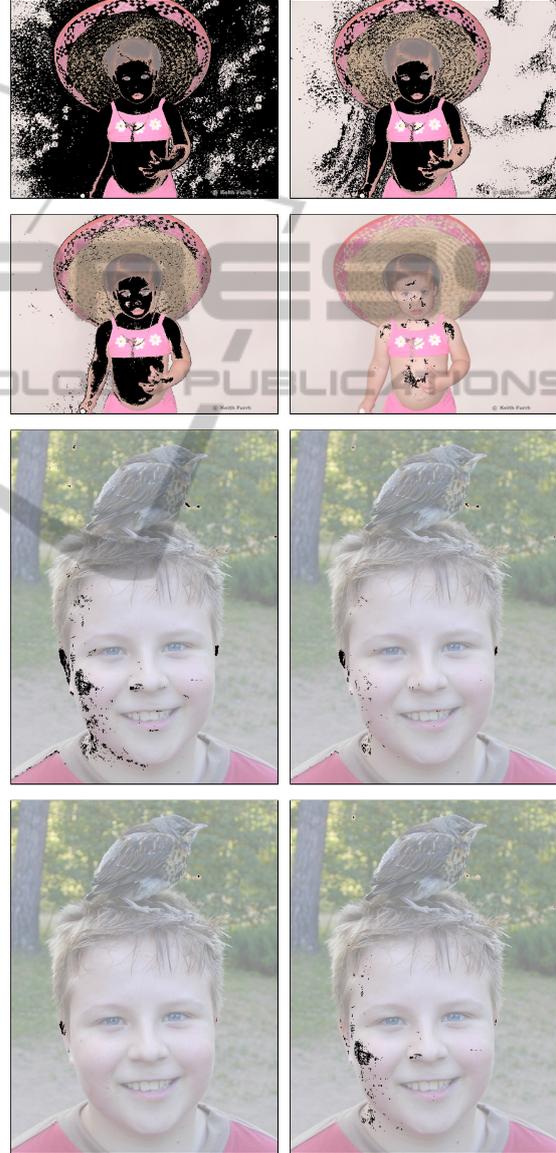
Furthermore, the seed detection process demonstrated to be sensitive to the bin size, performing better for more quantized color values. Another factor that influenced the seeds is the kernel size for the median filter. We empirically observed that the amount of seeds found is directly related to it. If it is too small, then very few seeds are found; otherwise, if it is too large, several false positives are placed as seeds. A  $15 \times 15$  median filter was used to generate the reported results.

A comparison between our seed extraction method and different fixed thresholds are presented in Table 2. Since an important issue in the seed extraction is to avoid false positives while retaining some true positives, the precision ( $\eta_{prec}$ ) seems appropriate for the evaluation. As it can be noticed, our method provides superior results with a large difference.

Another desired quality for seeds is that they should be spread over the skin regions to prevent from missing any isolated region. Therefore, we present a qualitative comparison in Figure 3. The seeds acquired with fixed thresholds are displayed along with the seeds collected by our method and the  $T_{seed}$  found

Table 2: Evaluation of seed extraction for validation data set.

Method	$\eta_{prec}$ (%)
$T_{seed} = 0.70$	68.32
$T_{seed} = 0.80$	73.39
$T_{seed} = 0.90$	82.64
$T_{seed} = 0.95$	88.70
Our seeds	<b>94.66</b>



(g)  $T_{seed} = 0.9$

(h) Our seeds [0.722982]

Figure 3: Examples of seeds detected through different thresholds for our method.

by it are placed in brackets.

It is possible to observe that our method not only avoids the non-skin regions but also maintains the dis-

persal. It is also noticeable that a high threshold is required in the first image in order to avoid misclassification, while the second will not produce significant seeds for the same value. Thus, our seed extraction method overcomes such problem with a choice of a proper threshold for each image.

The edge detection is an important stage of our method since it supports both seed extraction and propagation. Some experiments were conducted in the edge detector for different color spaces: only luminance, HSV, RGB and YCbCr. The HSV model better captured people's contours in images under abnormal lighting and, therefore, was employed in the experiments. Canny detector was applied with both low and high thresholds equal 100 and a dilation process was performed with a  $3 \times 3$  structuring kernel.

For comparison, we selected some state-of-the-art methods available in the literature: Cheddad's decision rule (Cheddad et al., 2009), statistical model (Jones and Rehg, 2002), face-based adaptation (Kawulok, 2010) built with Viola-Jones face-detector (Viola and Jones, 2004) and cost propagation (Kawulok, 2013). Our method was tested with  $\gamma = 0.8$  and  $T_{seed_{min}} = 0.5$ , whereas original parameters were employed in the other approaches.

Figure 4 presents a comparison of the ROC curves. The points in the curves were obtained with different thresholds, except for Cheddad's rule, whose output is binary.

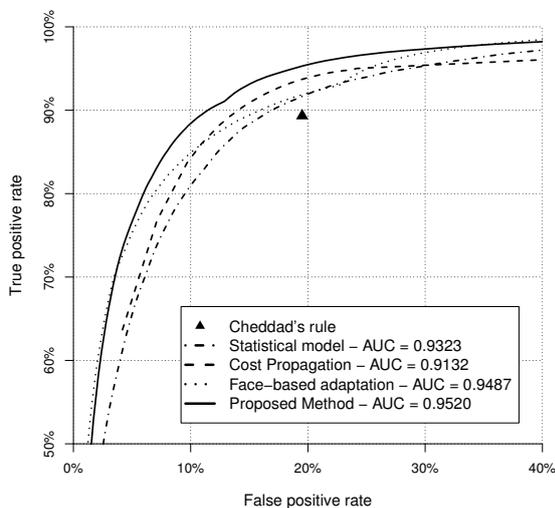


Figure 4: ROC curves for comparison of the tested methods.

To present quantitative values, Table 3 shows the results for a fixed threshold obtained through maximal Youden's index (Youden, 1950), which represents the closest point to the optimum value (0,1). Best values are highlighted in bold.

Cheddad's rule, as expected, holds the worst re-

Table 3: Detection results for different methods.

Method	$\eta_{tp}$ (%)	$\delta_{fp}$ (%)	$F_{score}$ (%)	$\delta_{min}$ (%)
Cheddad	89.33	19.51	64.78	30.18
Statistical model	87.90	14.51	69.71	26.61
Face-based adaptation	86.83	11.79	72.63	24.96
Cost propagation	<b>90.40</b>	14.46	71.05	24.06
Proposed method	89.78	<b>11.24</b>	<b>74.95</b>	<b>21.46</b>

sults, which demonstrates the inadequacy of such restrictive and biased method. Statistical model performs a little better, however, it still has a higher false positive rate.

From the compared methods, the cost propagation has the lowest  $\delta_{min}$  whereas the face-based adaptation the highest  $F_{score}$ . In fact, their values are very similar, considered as a tie. The first has  $\delta_{fp}$  very similar to the statistical model, what suggests that its improvement is accomplished by an increase in the true positive rate. Oppositely, the second maintains the  $\eta_{tp}$ , performing a reduction in the number of false positives. Our proposed method, which holds the best results, improves on both metrics.

Table 4 gives the true positive rate for fixed false positive rate values, where only methods of probabilistic output were considered. It shows the behavior of the methods in different tolerance settings.

Table 4: True positive rates for fixed values of false positive rate.

Method	$\eta_{tp}$ (%)	$\delta_{fp}$ (%)	$\eta_{tp}$ (%)	$\delta_{fp}$ (%)
Statistical model	81	10	92	20
Face-based adaptation	85	10	92	20
Cost propagation	84	10	94	20
Proposed method	<b>88</b>	10	<b>95</b>	20

Figure 5 illustrates the results for applying the evaluated methods on four different image samples from the test set. In the first row, it is possible to observe that the face-based adaptation misclassified a piece of blue shirt possibly due to the girl's blue eyes. Furthermore, in the second row shows the detection of several different false positives, which suggests that the face was not correctly detected.

The main drawback of the cost propagation approach is that it detects part of the background as skin, as illustrated in the third row of Figure 5. This occurs because, in some points, there is a smooth transition between skin and false skin regions, such that a "leakage" occurs.

Our method overcomes such problems since it tends to use more than just one region for the estimation. Thus, our local model is an accurate representation in cases of variation of skin through differ-



Figure 5: Examples of skin detected with different methods.

ent locations. Furthermore, seeds were always found in our tests, while no faces were detected in 12% of the images. Although “leakages” can still occur, they are significantly reduced as demonstrated through the results.

It is also important to highlight the viability of our method in real-time applications, since the average time per image in the test set was 282ms in an un-optimized version of our code.

## 5 CONCLUSIONS AND FUTURE WORK

This work presented a new adaptive human skin segmentation method that makes use of skin probability map and eliminates the need for object detection. The main contributions of our approach include: a new method for seed extraction based on spatial analysis and a self-contained adaptation.

Experimental results demonstrated that the proposed technique outperforms state-of-the-art skin segmentation methods for a large and well-known test set. Nevertheless, additional improvements can be made. False positives generated in the propagation

stage of our method has a large contribution to the overall accuracy, what makes us conjecture that even little enhancements in the propagation control, such as better edge detection, will significantly decrease the error rates.

Our method could also be combined with face-based adaptive methods through two strategies: use of detected faces to improve seed extraction and use as an alternative when a face is not present in the image or has not been correctly detected.

As future directions, we intend to investigate more powerful features, such as textural information, to discard incorrect seeds, as well as new strategies for controlling the propagation.

## ACKNOWLEDGEMENTS

The authors are grateful to FAPESP, CNPq and CAPES for the financial support.

## REFERENCES

Cheddad, A., Condell, J., Curran, K., and Mc Kevitt, P. (2009). A Skin Tone Detection Algorithm for an

- Adaptive Approach to Steganography. *Signal Processing*, 89(12):2465–2478.
- Dijkstra, E. W. (1959). A Note on Two Problems in Connection with Graphs. *Numerische Mathematik*, 1:269–271.
- Fleck, M. M., Forsyth, D. A., and Bregler, C. (1996). Finding Naked People. In *European Conference on Computer Vision*, pages 593–602. Springer.
- Fritsch, J., Lang, S., Kleinhagenbrock, M., Fink, G. A., and Sagerer, G. (2002). Improving Adaptive Skin Color Segmentation by Incorporating Results from Face Detection. In *11th IEEE International Workshop on Robot and Human Interactive Communication*, pages 337–343.
- Gonzalez, R., Woods, R., and Eddins, S. (2009). *Digital Image Processing Using MATLAB*. Gatesmark Publishing.
- Hsu, R.-L., Abdel-Mottaleb, M., and Jain, A. K. (2002). Face Detection in Color Images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):696–706.
- Jones, M. J. and Rehg, J. M. (2002). Statistical Color Models with Application to Skin Detection. *International Journal of Computer Vision*, 46(1):81–96.
- Kakumanu, P., Makrogiannis, S., and Bourbakis, N. (2007). A Survey of Skin-Color Modeling and Detection Methods. *Pattern Recognition*, 40(3):1106–1122.
- Kawulok, M. (2010). Energy-based Blob Analysis for Improving Precision of Skin Segmentation. *Multimedia Tools and Applications*, 49(3):463–481.
- Kawulok, M. (2013). Fast Propagation-based Skin Regions Segmentation in Color Images. In *10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition*, pages 1–7.
- Kawulok, M., Nalepa, J., and Kawulok, J. (2014). Skin Detection and Segmentation in Color Images. In *Advances in Low-Level Color Image Processing*, pages 329–366. Springer.
- Kovac, J., Peer, P., and Solina, F. (2003). Human Skin Color Clustering for Face Detection. In *EUROCON - Computer as a Tool*, volume 2, pages 144–148. IEEE.
- Lee, J.-Y. and Yoo, S. I. (2002). An Elliptical Boundary model for Skin Color Detection. In *International Conference on Imaging Science, Systems, and Technology*. Citeseer.
- Phung, S. L., Bouzerdoum, A., and Chai Sr, D. (2005). Skin segmentation using color pixel classification: Analysis and comparison. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(1):148–154.
- Phung, S. L., Chai, D., and Bouzerdoum, A. (2003). Adaptive Skin Segmentation in Color Images. In *International Conference on Multimedia and Expo*, volume 3, pages III–173.
- Ruiz-del Solar, J. and Verschae, R. (2004). Skin Detection using Neighborhood Information. In *Sixth IEEE International Conference on Automatic Face and Gesture Recognition*, pages 463–468. IEEE.
- Sobotka, K. and Pitas, I. (1998). A Novel Method for Automatic Face Segmentation, Facial Feature Sxtraction and Tracking. *Signal Processing: Image Communication*, 12(3):263–281.
- Soriano, M., Martinkauppi, B., Huovinen, S., and Laaksonen, M. (2000). Skin Detection in Video under Changing Illumination Conditions. In *15th International Conference on Pattern Recognition*, volume 1, pages 839–842. IEEE.
- Subban, R. and Mishra, R. (2014). Human skin segmentation in color images using gaussian color model. In *Recent Advances in Intelligent Informatics*, pages 13–21. Springer.
- Taylor, M. J. and Morris, T. (2014). Adaptive Skin Segmentation via Feature-based Face Detection. In *SPIE Photonics Europe*, pages 91390P–91390P. International Society for Optics and Photonics.
- Vezhnevets, V., Sazonov, V., and Andreeva, A. (2003). A Survey on Pixel-based Skin Color Detection Techniques. In *Graphicon*, volume 3, pages 85–92. Moscow, Russia.
- Viola, P. and Jones, M. J. (2004). Robust Real-Time Face Detection. *International Journal of Computer Vision*, 57(2):137–154.
- Wang, Y. and Yuan, B. (2001). A Novel Approach for Human Face Detection from Color Images under Complex Background. *Pattern Recognition*, 34(10):1983–1992.
- Yang, M.-H. and Ahuja, N. (1999). Gaussian Mixture Model for Human Skin Color and its Application in Image and Video Databases. In *SPIE: Storage and Retrieval for Image and Video Databases VII*, volume 3656, pages 458–466.
- Youden, W. J. (1950). Index for Rating Diagnostic Tests. *Cancer*, 3(1):32–35.