# FAST ADAPTABLE SKIN COLOUR DETECTION IN RGB SPACE

Martin Tosas[1] and Steven Mills[2]

[1] *School of Computer Science and IT, University of Nottingham, Nottingham, United Kingdom*

[2] *Geospatial Research Center (NZ) Ltd, University of Canterbury, Christchurch, New Zealand*

Keywords: Skin Colour Detection, Linear Container, Human Computer Interaction.

Abstract: This paper presents a skin colour classifier that uses a linear container in order to confine a volume of the RGB space where skin colour is likely to appear. The container can be adapted, using a single training image, to maximize the detection of a particular skin tonality. The classifier has minimum storage requirements, it is very fast to evaluate, and despite operating in the RGB space, provides equivalent illumination (brightness) independence to that of classifiers that work in the rg-plane. The performance of the proposed classifier is evaluated and compared with other classifiers. Finally, conclusions are drawn.

## 1 INTRODUCTION

We propose a skin colour classifier named the Linear Container (LC) classifier. The classifier uses four decision planes in order to confine a volume of the RGB space where skin colours are likely to appear. The features of the LC classifier are: capable of being tuned to a particular skin tonality; rapid evaluation; minimal storage requirements; and resistance to illumination (brightness) changes equivalent to that of classifiers that work in normalised RGB. The classifier requires a tuning stage in which a single image, with marked skin and background areas, is analysed resulting in a new linear container configuration. The classifier is originally intended to detect skin colour in tracking applications, such as the ones found in Human Computer Interaction (HCI) systems; however, other uses are conceivable.

The paper is organized as follows: Section 2 outlines previous work on skin colour detection, and situates the LC classifier in relation to other works; Section 3 introduces the LC classifier, and describes its tuning procedure; Section 4 presents performance results; Section 5 shows the behaviour of the LC classifier when tuned at various resolutions; Section 6 studies some HCI usability factors; Finally, Section 7 gives some conclusions and directions for future work.

## 2 PREVIOUS WORK ON SKIN COLOUR DETECTION

Skin colour provides an important source of information for computer vision systems that monitor people. The skin colour cue is widely used in face detection and recognition systems, various types of surveillance, vision-based biometric systems, and vision-based HCI systems. All these areas of application use skin colour to track, locate and interpret people, with relatively efficient, fast, low-level, methods.

The goal of skin colour detection is to build a decision rule that can discriminate between the skin and non-skin colour pixels of an image. Because of the importance of skin colour detection there have been numerous approaches to solve this task. The various approaches can be grouped into the following four categories: non-parametric skin distribution modelling, parametric skin distribution modelling, explicitly defined skin region modelling, and dynamic skin colour modelling (Vezhnevets et al., 2003).

Non-parametric skin distribution modelling uses training data to estimate a skin colour distribution. This estimation process is sometimes referred to as the construction of a Skin Probability Map (SPM) (Jones and Rehg, 1999; Brand and Mason, 2000; Gomez, 2002) assigning a probability value to each



Tosas M. and Mills S. (2007).
FAST ADAPTABLE SKIN COLOUR DETECTION IN RGB SPACE.
In *Proceedings of the Second International Conference on Computer Vision Theory and Applications*, pages 3-10
DOI: 10.5220/0002055800030010

point of a discretised colour space. A SPM can be implemented by a colour histogram, and such approaches normally use the chrominance plane of some colour space in order to offer resistance to illumination changes (Chen et al., 1995; Schumeyer and Barner, 1998; Jones and Regh, 1999; Zarit et al., 1999). SPMs can use a Bayes classification rule in order to improve their performance, in this case two colour histograms are required; one for the probability of skin colour, and another for the probability of non-skin colour (Jones and Regh, 1999; Zarit et al., 1999; Chai and Bouzerdoum, 2000). The main disadvantages of SPMs are the high storage requirements and the fact that their performance directly depends on the representativeness of the training images.

Parametric skin distribution modelling can represent skin colour in a more compact form. Common examples of parametric modelling model a skin colour distribution using a single Gaussian (Ahlberg, 1999; Menser and Wien, 2000; Terrillon et al., 2000) or a mixture of Gaussians (Jones and Regh, 1999; Yang and Ahuja, 1999; Terrillon et al., 2000). Expectation Maximization (EM) algorithms are used on training data to find the model parameters that produce the best fit. The goodness of fit, and therefore the performance of the model, depends on the shape of the chosen model and the chosen colour space. This performance dependency with the colour space is stronger in the case of parametric modelling than it is in the case of non-parametric modelling (Brand and Mason, 2000; Lee and Yoo, 2002).

Another way to build a skin colour classifier is to define explicitly, through a number of rules, the boundaries of a skin cluster in some colour space; this is called explicitly defined region modelling. The obvious advantage of this method is its computational simplicity, which has attracted many researchers (Fleck et al., 1996; Ahlberg, 1999; Jorda et al., 1999; Peer et al., 2003), as it leads to the construction of a very rapid classifier. However in order to achieve high recognition rates both a suitable colour space and adequate decision rules need to be found empirically. Gomez and Morales (2002) proposed a method that can build a set of rules automatically by using machine learning algorithms on training data. They reported results comparable to the Bayes SPM classifier in RGB space for their data set.

Finally, we have dynamic skin colour modelling. This category of skin modelling methods is designed for skin detection during tracking. Skin detection in this category is different from static image analysis in a number of aspects. First, in principle, the skin models in this category can be less general – i.e tuned for a specific person, camera, or lighting. Second, an initialisation stage is possible, when the skin region of interest is segmented from the background by a different classifier or manually; this makes possible to obtain a skin classification model that is optimal for the given conditions. Finally, this category of skin models can be able to update themselves in order to match changes in lighting conditions. Some of the methods in this category use Gaussian distribution adaptation (Yang and Ahuja, 1998), or dynamic histograms (Soriano et al., 2000; Stern and Efros, 2002). In (Soriano et al., 2000) a skin locus, in rg space, is constructed beforehand from training data. Then, during tracking, their dynamic skin colour histogram is updated with pixels from the bounding box of the target, provided these pixels belong to the skin locus. This makes the dynamic histogram less likely to adapt to colour distributions other than that of skin.

The proposed LC classifier belongs to the last two categories. The classifier is implemented using rules similar to the rules of the explicitly defined skin region models; however, these rules are parameterised in order that they can be tuned to specific conditions, during an initialisation stage. The parameters of the LC classifier can also be recalculated rapidly in order to adapt to changing illumination conditions.

# 3 LINEAR CONTAINER CLASSIFIER

Normalised RGB is a popular colourspace because of its simple normalisation procedure, and its diminished dependence with brightness (Yang and Ahuja, 1998; Zarit et al., 1999; Lee and Yoo, 2002; Stern and Efros, 2002; Peer et al., 2003). The projection from RGB (3D space) to normalised RGB (2D space) corresponds with a cone in the original 3D RGB space, in that each point in the rg-plane corresponds to a 3D line of colour values in the original RGB space. These lines meet at (0, 0, 0), and points along the lines correspond to scaling of white illumination. Therefore, a skin colour cluster in the rg-plane corresponds to a cone-like cluster in RGB space. This is illustrated in Figure 1.

The proposed LC classifier uses a polyhedral cone, constructed from four decision planes, in order to model the cone-like region in RGB space that

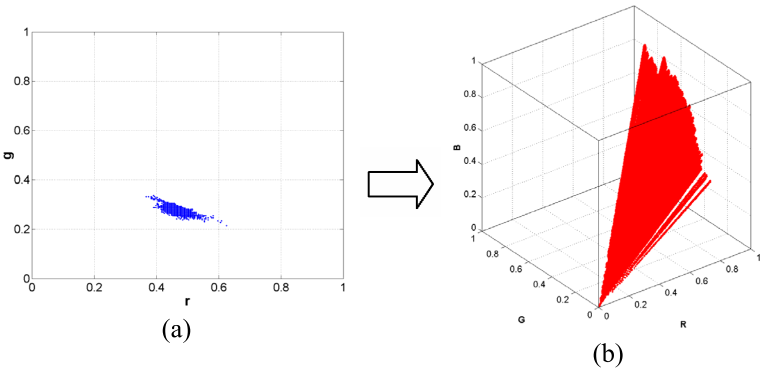


Figure 1: (a) Skin colour cluster, in the rg-plane, from a single sample. (b) The rg-plane skin colour cluster projected to RGB space; each point in the rg-plane becomes a line in RGB space.

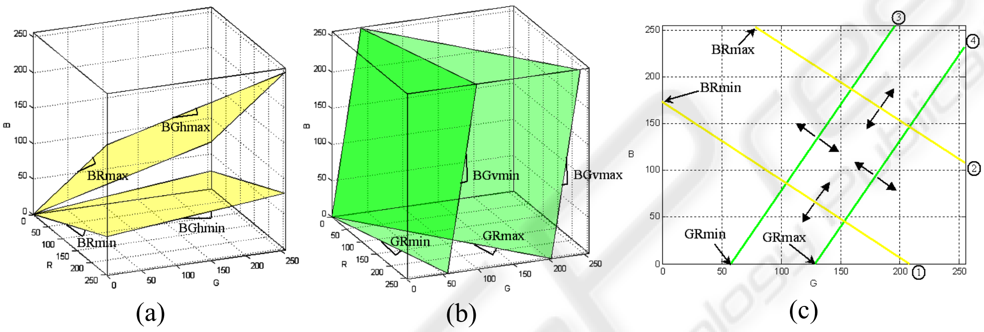


Figure 2: (a) Horizontal decision planes. (b) Vertical decision planes. (c) Tuning heuristic.

results from the projection of a skin colour cluster in the rg-plane to the RGB space. The LC classifier performs pixel-based segmentation. If an RGB value is inside the polyhedral cone volume, it is classified as skin; if the RGB value is outside the polyhedral cone volume then it is classified as non-skin. The definition of the four decision planes is:

$$BGhmin \cdot G + BRmin \cdot R < B < BGhmax \cdot G + BRmax \cdot R \quad (1)$$

where BGhmin and BRmin parameterise the lower "horizontal" plane, and BGhmax and BRmax parameterise the higher horizontal plane. The horizontal planes are illustrated in Figure 2(a). These two planes confine a volume between them by constraining the values that B can take in relation to R and G. This volume is further constrained by two "vertical" planes:

$$BGvmin \cdot B + GRmin \cdot R < G < BGvmax \cdot B + GRmax \cdot R \quad (2)$$

where BGvmin and GRmin parameterise the left vertical plane, and BGvmax and GRmax parameterise the right vertical plane. The vertical planes confine a volume between them by constraining the values that G can take in relation to R and B. The vertical planes are illustrated in Figure2(b). As the RGB values that are close to the origin carry too little colour information, we truncate the apex of the polyhedral cone with one additional rule: $Rmin < R$. If a colour value satisfies these three rules, then it is inside the polyhedral cone, and therefore classified as skin colour.

The LC classifier can be tuned for a specific person, camera, or lighting conditions, in an initialisation step. For this, an initialisation image is needed. This initialisation image is composed of two approximately complementary masks; one mask delimits the target skin colour area, we call this mask *SkinMask*; and the other mask comprises areas where we do not expect to find skin colour, we call this mask *BackgroundMask*. Figure 3 shows an initialisation image segmented by the two masks. The BackgroundMask can be tailored in order to avoid areas of skin colour in addition to those included in SkinMask, for example, Figure 3(b) avoids the subject's wrist. The two masks can be generated manually, or automatically by a tracking system such as (Tosas and Li, 2007).

The tuning procedure uses a heuristic method by which the parameters of the decision planes are changed in sequence.

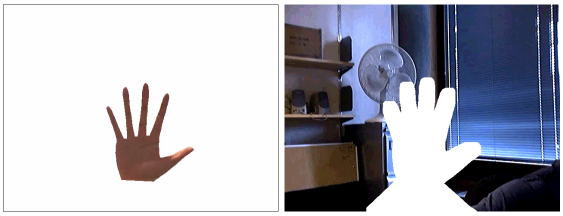

(a) (b)

Figure 3: Initialisation image segmented by SkinMask (a) and BackgroundMask (b).

Each time a parameter is changed, the *fitness* of the LC classifier, to the detection of skin colour in the SkinMask and to the rejection of skin colour in the BackgroundMask, is measured using the following equation:

$$fitness = \frac{\#\, skin\ in\ SkinMask}{size\ of\ SkinMask} \times TI - \frac{\#\, skin\ in\ BackgroundMask}{size\ of\ BackgroundMask} \quad (3)$$

where *TI* (Target Importance) is used to control the importance of the target skin colour area in the fitness. In the experiments of the following sections *TI = 2* so as to give double importance to detecting skin on the SkinMask than to avoid detecting skin on the BackgroundMask. This parameter allows the classifier to be tuned to favour true positives or negatives.

The heuristic search, by which the parameters of the decision planes are changed, is illustrated in Figure 2(c). This figure shows a section view of the RGB cube, corresponding to the B-G-plane with maximum R. Lines 1, 2, 3 and 4 are the intersections of the four decision planes with the section view. Starting from some priori values, the search varies BRmin, then BRmax, GRmin, and finally GRmax; first, reducing their values, then increasing their values, and measuring the fitness (Equation 3) at each step. The values that produce the best fitness are finally selected. Note that the slope of the decision planes remains unchanged in this heuristic.

## 4 PERFORMANCE RESULTS

The LC classifier is tested on video sequences of subjects with four different skin tonalities: Mediterranean, white Caucasian, black African, and Chinese. The target skin colour area is the subject's hand. The subjects hold their hand open in front of the camera, and move the hand towards and away from the camera. An overhead lamp affects the illumination of the subjects' hand. When the subject's hand is closer to the camera, the hand is under a shadow and looks darker. When the subject's hand is further away from the camera, the hand is under the lamp and looks brighter. The classifier is initialised once, at the beginning of each video sequence.

The skin colour detection performance is calculated for each video sequence, using a ground truth. The ground truth consists of two masks, which have been manually generated for every fifth frame of the four video sequences. The ground truth considers the subject's hand as the target area for skin colour detection. This area is segmented using the *SkinTruth* mask, Figure4(b). The background is segmented using the *BackgroundTruth* mask, Figure 4(c). Note that the BackgroundTruth mask is not the complement of the SkinTruth mask. The BackgroundTruth mask avoids the target skin colour area, the subject's hand, and any other skin colour areas in the image; therefore, for each measurement frame, there will be some areas which will not take part in the counting; these areas correspond to the subject's face and arms. Both masks are tested for skin colour. Skin colour pixels found in the SkinTruth mask constitute true-positives. Non-skin colour pixels found in the BackgroundTruth mask constitute true-negatives. In order to compare detection results between frames the true-positives and true-negatives are normalised to the size of SkinTruth and BackgroundTruth masks respectively. Normalised true-positives are referred to as NTN, and normalised true-negatives are referred to as NTP.

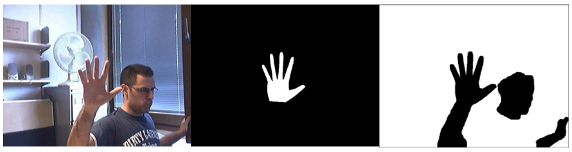

Figure 4: (a) Original frame. (b) SkinTruth mask. (c) BackgroundTruth mask.

We use an rg skin colour histogram classifier as a comparison reference with the LC classifier. The rg histogram used for comparison is constructed in an initialisation step at the beginning of each sequence from the pixels in SkinMask and its size is 64x64 bins. A pixel is classified as skin colour if its corresponding bin in the rg histogram is bigger than a threshold.

The choice of the threshold affects the detection rate of the rg histogram. In general, if the threshold increases, NTN tends to be higher, but NTP tends to be lower; if the threshold decreases, NTP tends to be higher, but NTN tends to be lower. For the tested video sequences a threshold of 25 produced the best results.

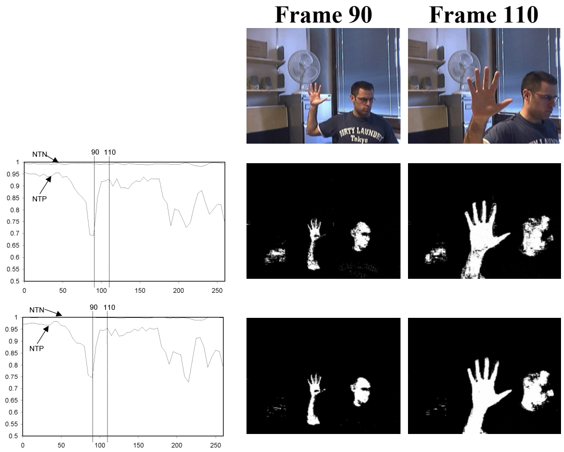


Figure 5: Mediterranean subject test. Top row: Original frames. Middle row: rg histogram classifier. Bottom row: LC classifier.

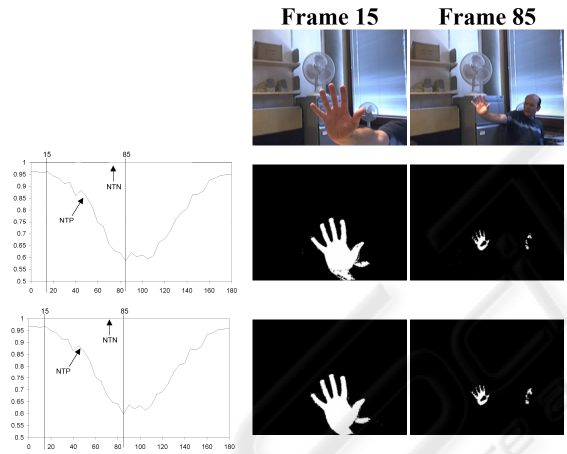


Figure 7: White Caucasian subject test. Top row: Original frames. Middle row: rg histogram classifier. Bottom row: LC classifier.

Figure 5 shows the results for the Mediterranean subject. The figure presents plots of the NTP and NTN against the frame number, and two example frames showing the skin colour classification for a best detection case and a worst detection case. The top row shows the original frames before detecting the skin colour areas. The middle row corresponds to the rg histogram, and the bottom row corresponds to the LC classifier. The results of both classifiers are similar, but the LC classifier consistently exhibits slightly larger NTP and NTN than the rg histogram classifier, all along the video sequence. Following the same format as Figure 5, Figures 6, 7 and 8 show the results for the other three ethnic skin tonalities. In all cases, the LC classifier exhibited the same or larger NTP and NTN than the rg histogram classifier.

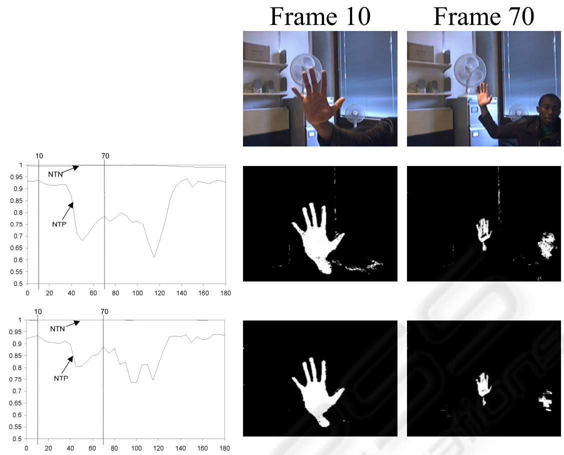


Figure 6: Black African subject test. Top row: Original frames. Middle row: rg histogram classifier. Bottom row: LC classifier.

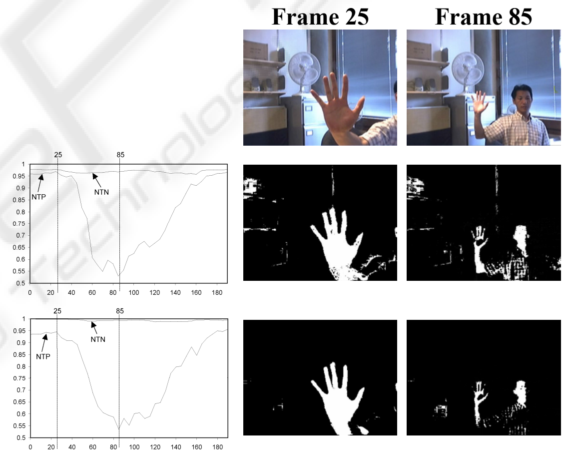


Figure 8: Chinese subject test. Top row: Original frames. Middle row: rg histogram classifier. Bottom row: LC classifier.

Videos showing the skin colour detection tests are available at: `www.cs.nott.ac.uk/~mtb/research/SkinColour.html`

An experiment comparing the computational speed of the LC classifier, against other classifiers, was also carried out. The experiment consists in measuring the time it takes for a classifier to check all the pixels in a 640x480 frame. The experiment is repeated for 100 frames of a video sequence containing skin colour regions, and the times used in each frame are averaged.

Table 1: Execution time results.

| | Average time per frame | Speed-up of the RGB LC classifier with respect to the other classifiers |
|---|---|---|
| RGB LC classifier | 0.0090 secs | |
| rg LC classifier | 0.0147 secs | x1.62 |
| rg LC classifier with lookup table | 0.0107 secs | x1.17 |
| rg histogram | 0.0235 secs | x2.59 |
| rg histogram with lookup table | 0.0204 secs | x2.25 |
| bare RGB histogram | 0.0022 secs | x0.24 |

The experiment was carried out in an AMD Athlon 3500+, 1GB of RAM. The results are shown in Table 1.

The RGB LC classifier uses an extra rule to avoid dark pixels, the other classifiers do not use this rule. The rg LC classifier is the 2D equivalent to the proposed RGB LC classifier. It works in the rg-plane by using 4 decision lines instead of 4 decision planes. The skin detection performance of this classifier is equivalent to the RGB LC classifier. The equations in the rg LC classifier are simpler than those of the RGB LC classifier; however, the former is slower because it has to normalise each pixel from RGB to rg. The use of lookup table containing all the possible normalisations can speed up the normalisation procedure. But, even when using a lookup table, the RGB LC classifier is x1.172126 times faster than its rg LC equivalent. And the rg LC classifier is faster than the rg histogram classifier (both with and without lookup table).

A bare RGB histogram classifier is used as a comparison measure for the fastest skin colour classifier (yet the most sensitive to changes in illumination), being x4.16 times faster than the LC classifier. However, in practice, the reduced storage requirements of the LC classifier may result, when plugging it into certain algorithms, in faster execution times due to its better locality of reference. This is illustrated in a practical application of the LC skin colour classifier. The LC classifier is used in the measurement function of the particle-filter based hand contour tracking algorithm described in (Tosas and Li, 2007). In such a tracking algorithm, most of the computation is expended in the measurement function (profiling shows that more than 60% of the application's time is spent in the measurement function). The average time spent in the processing of a tracking time-step is calculated as the average of the times spent in each of 100 time-steps of tracking. When exchanging the LC classifier for a 32x32x32 bins RGB histogram (in the same machine as the previous experiment) the speed up in the processing of a time-step is only x1.2 times faster (as opposed to the x4.16 times faster suggested in the previous experiment).

## 5 TUNING AT VARIOUS RESOLUTIONS

So far, the LC classifier has been tuned using an initialisation image of the same size as the video sequence in which it was tested, 640x480 pixels. It was observed that the tuning of the LC classifier on a decimated version of the initialisation image, results in little degradation of the classifier's detection performance on the non-decimated video sequence. This is because the result of the tuning is more dependent upon the range of colours of the pixels in the initialisation masks than upon the number of pixels. This fact allows us to speed-up the tuning procedure, because the amount of data to be dealt with is reduced. The speed-up of the tuning procedure as a result of using a decimated initialisation image instead of using a non-decimated initialisation image is: x4 for a 320x240 resolution, x16 for 160x120, x64 for 80x60, x256 for 40x30, and x1024 for a 20x15 resolution. Figure 9 shows the NTP of the LC classifier, on the video sequence of the Mediterranean subject, for various resolutions

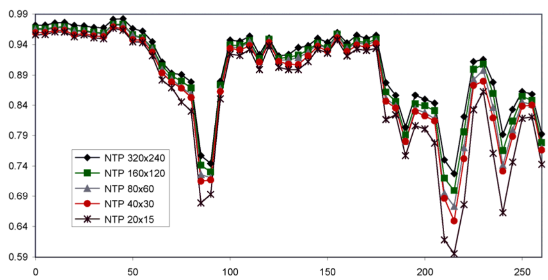


Figure 9: NTP when tuning at various resolutions.

of the initialisation image. The NTN are not shown as they remain at almost 1 for the six resolutions. Notice that the NTP for an initialisation image of 320x240 is virtually the same as the NTP for an initialisation image of 640x240.

The speed-up resulting from the use of decimated initialisation images becomes very important in HCI applications because it allows the tuning (and potentially periodical retuning) of the LC classifier with only a small impact in the HCI system responsiveness.

# 6 HCI USABILITY FACTORS

The tuning stage in the experiments of the previous sections was idealised, in that no background colours appear in the SkinMask, and no skin colour appeared in BackgroundMask. If the LC classifier is used in a HCI system, which could generate the initialisation masks automatically from a tracking subsystem, it is possible that background appears in SkinMask, and skin colour appears in BackgroundMask. In this section we study the robustness of the LC classifier against non-ideal tuning conditions.

The detection performance of the LC classifier is calculated, once more, for the video sequence of the Mediterranean subject. This time, the tuning is repeated for a misaligned SkinMask and BackgroundMask. In each repetition SkinMask only contains a percentage of the target's skin colour area. The skin that is not in the SkinMask is in the BackgroundMask, this affects the final configuration of the LC parameters found during the tuning stage. Figure 10 shows the NTP for four percentages of skin colour in SkinMask. The NTN is not shown as it is almost unaffected in all the four cases. We can see that the degradation in NTP for a 50% skin in SkinMask is small; and even when the amount of skin in SkinMask is as small as 25%, the NTP along the whole sequence may still be useful for some applications. However, the model parameters found during the tuning stage, depend on the colours appearing in each initialisation mask; hence, different results are possible even when SkinMask contains the same amount of skin. This is illustrated in Figure 11, where the tuning of the LC classifier using two SkinMasks with the same percentage of skin inside the mask, produce different detection performances.

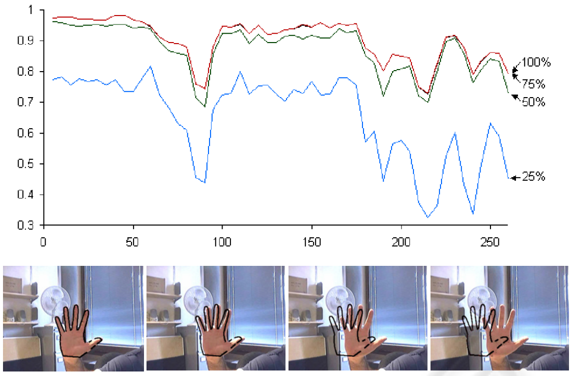


(a) (b) (c) (d)

Figure 10: Chart, NTP for four percentages of skin in SkinMask. (a) SkinMask containing 100% skin, (b) 75% skin, (c) 50% skin, and (d) 25% skin.

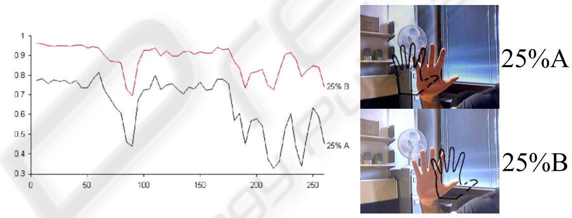


Figure 11: NTP for two SkinMask containing 25% skin colour.

# 7 CONCLUSIONS

We have presented the Linear Container (LC) skin colour classifier. This classifier constitutes a contribution to dynamic skin colour modelling methods. Its detection performance compares well with an rg histogram classifier, resulting in equal or better detection rates, when using a single training image. Two remarkable qualities of this classifier are its evaluation speed, and its low storage requirements. The four rules that define the decision planes, and an extra rule to avoid dark pixels, can be rapidly evaluated, resulting in a x2.24 speed-up with respect to a simple rg histogram classifier. In practice, the reduced storage requirements of the LC classifier may result, when plugging it into certain algorithms, in even faster execution times due to its better locality of reference. This can prove to be an advantage in embedded systems. On the other hand, despite the LC classifier operates in the RGB space, its resistance to illumination changes is equivalent to that of a classifier that operates in the rg-plane. The detection performance of the LC classifier is not greatly impaired when the tuning is performed in a decimated initialisation image, but the execution time of the tuning is notably reduced. The LC classifier also proved to be robust to non-ideal initialisations, in which skin colour appears in BackgroundMask, and background appears in SkinMask.

A subject of further work is the tuning stage. Different heuristics or maximisation procedures could produce better detection results. On the other hand, the LC model itself could be changed. Linear containers are fast to evaluate, but other type of containers, could produce a better fit of the skin colour cluster through scaling of white illumination. Sets of rules such as the ones proposed in (Gomez and Morales, 2002) could give better detection results, although a tuning procedure for these rules may be more complex.

## REFERENCES

Ahlberg, J., 1999. A system for face localization and facial feature extraction. Tech. Rep. LiTH-ISY-R-2172, Linkoping University.

Brand, J. and Mason, J., 2000. A comparative assessment of three approaches to pixel level human skin-detection. In *Proc. of the ICPR*, vol. 1, 1056-1059.

Chai, D., And Bouzerdoum, A., 2000. A Bayesian approach to skin color classification in ycbcr color space. In *Proc. of IEEE Region Ten Conference (TENCON'2000)*, vol. 2, 421-424.

Chen, Q., Wu, H., and Yachida, M., 1995. Face detection by fuzzy pattern matching. In *Proc. of the ICCV*, 591-597.

Fleck, M., Forsyth, D. A., and Bregler, C., 1996. Finding naked people. In *Proc. of the ECCV*, vol. 2, 592-602.

Gomez, G., 2002. On selecting colour components for skin detection. In *Proc. of the ICPR*, vol. 2, 961-964.

Gomez, G., and Morales, E., 2002. 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.

Jones, M. J. and Rehg, J. M., 1999. Statistical color models with application to skin detection. In *Proc. of the CVPR '99*, vol. 1, 274-280.

Jorda, L., Perrone, M., Costeira, J., and Santos-Victor, J., 1999. Active face and feature tracking. In *Proc. of the 10th International Conference on Image Analysis and Processing*, 572-577.

Lee, J. Y., and Yoo, S. I., 2002. An elliptical boundary model for skin color detection. In *Proc. of the 2002 International Conference on Imaging Science, Systems, and Technology*.

Menser, B., and Wien, M., 2000. Segmentation and tracking of facial regions in color image sequences. In *Proc. Visual Communications and Image Processing, SPIE*, 731-740.

Peer, P., Kovac, J., and Solina, F., 2003. Human skin colour clustering for face detection. In *International Conference on Computer as a Tool*, EUROCON, The IEEE region 8, vol 2, 144-148.

Vezhnevets, V., Sazonov V., Andreeva A., 2003. A Survey on Pixel-Based Skin Color Detection Techniques. In *Proc. Graphicon*, pp. 85-92, Moscow, Russia.

Schumeyer, R., And Barner, K., 1998. A color-based classifier for region identification in video. In *Proc.Visual Communications and Image Processing, SPIE*, vol 3309, 189-200.

Soriano, M., Martinkauppi, B., Huovinen, S., and Laaksonen, M., 2000. Using the skin locus to cope with changing illumination conditions in color-based face tracking. In *Proc. of the IEEE Nordic Signal Processing Symposium*, pp. 383-386.

Stern, H., and Efros, B., 2002. Adaptive color space switching for face tracking in multi-colored lighting environments. In *Proc. of the International Conference on Automatic Face and Gesture Recognition*, 249-255.

Terrillon, J. C., Shirazi, M. N., Fukamachi, H., and Akamatsu, S., 2000. 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*, 54-61.

Tosas, M., and Li, B., 2007. Tracking Tree-Structured Articulated Objects Using Particle Interpolation. Accepted for publication in the *proceedings of CGIM2007*.

Yang, M., and Ahuja, N., 1998. Detecting human faces in color images. In *Proc. of ICIP*, vol. 1, 127-130.

Yang, M. H., and Ahuja, N., 1999. Gaussian mixture model for human skin color and its applications in image and video databases. In *Proc. of the SPIE: Conf. On Storage and Retrieval for Image and Video Databases*, vol. 3656, 458-466.

Zarit, B. D., Super, B. J. and Quek, F. K. H., 1999. Comparison of five color models in skin pixel classification. In *ICCV'99 Int'l Workshop on recognition, analysis and tracking of faces and gestures in Real-Time systems*, 58- 63.