

# Joint Adaptive Colour Modelling and Skin, Hair and Clothing Segmentation Using Coherent Probabilistic Index Maps

Carl Scheffler  
carl.scheffler@gmail.com

Jean-Marc Odobez  
odobez@idiap.ch

Idiap Research Institute  
Centre du Parc  
Rue Marconi 19  
Martigny, Switzerland

---

## Abstract

We address the joint segmentation of regions around faces into different classes — skin, hair, clothing and background — and the learning of their respective colour models. To this end, we adopt a Bayesian framework with two main elements. First, the modelling and learning of prior beliefs over class colour distributions of different kinds, *viz.* discrete distributions for background and clothing, and continuous distributions for skin and hair. This component of the model allows the formulation of the colour adaptation task as the computation of the posterior beliefs over class colour distributions given the prior (general) class colour model, a colour likelihood function and observed pixel colours. Second, a spatial prior based on probabilistic index maps enhanced with Markov random field regularisation. The inference of the segmentation and of the adapted colour models is solved using a variational scheme. Segmentation results on a small database of annotated images demonstrate the impact of the different modelling choices.

## 1 Introduction

We address the joint segmentation of regions around faces into different classes (skin, hair, clothing, and background) and the learning of their respective colour models. Such colour information is known to be a powerful cue for tracking, as has been demonstrated, for instance, using the mean shift [3] and particle filtering [14] algorithms. In particular, locating skin-coloured regions plays an important role when dealing with faces, and is useful not only for face detection and tracking [21, 22], but also for head pose estimation when precisely segmented and when dealing with low to mid-resolution images. Furthermore, knowing the colour of a person's skin, hair and clothing can aid in re-associating lost tracks in video.

Much of the existing literature on skin colour modelling is about building a *general* colour model across all possible ethnicities and illumination conditions. Such a model can still be distracted by objects that are approximately skin-toned, like wood and T-shirts. We aim to improve skin colour modelling and segmentation by using three main elements. The first is automatic adaptation of the colour model, either by updating some model parameters (*e.g.* of a Gaussian), or by adapting a threshold, or by doing both. The second is to perform colour segmentation, which brings two main benefits: a form of spatial regularisation and,

importantly, the competition between several colour models for class assignment, which resolves the issue of setting a threshold for the skin colour model. These two elements — adaptation and segmentation — are obviously coupled. The main challenge is to obtain accurate colour models, which depends on the choice of good priors for these models and the selection of appropriate data (pixels) for their estimation or adaptation. The third element can mitigate this data selection issue. It consists of the exploitation of a spatial prior on the location of each class label. The intuition is that knowing the approximate location of, for instance, a person’s hair and clothing provides a lot of information about where his/her face is located (and *vice versa*). A spatial prior also helps to avoid drift, explicitly for segmentation and implicitly for colour adaptation. In this paper, we combine the three elements above within a generative model of face images. More precisely, the main components and novelties of our model are as follows.

- The colour model of each class (skin, hair, clothing and background) is represented with a probability distribution over colours, which we will refer to as a *palette*. For each class, we learn a prior distribution over palettes — *i.e.*, a distribution over colour distributions — across all people. This contrasts with the literature where a prior model is an expected colour distribution across all people. The parameters of this prior can then be updated online, to do adaptation. Furthermore, we use different types of colour models for different classes — continuous models for skin and hair, and discrete models for background and clothing.
- We use a spatial prior over classes based on probabilistic index maps (PIMs) [9]. The standard PIM model is modified by incorporating a Markov random field prior to model correlations between neighbouring pixels.
- Inference in the full Bayesian model is done using a variational approach that computes approximate Bayesian posteriors over the class of each image pixel for segmentation and over the palette of each class for colour adaptation.

The impact and validity of the different components are evaluated on a small dataset.

The remainder of the paper is organised as follows. Section 2 discusses related work on skin and hair colour modelling. The proposed Bayesian model is described in Section 3. The data set used for training and testing is presented in Section 4. The different inference procedures are summarised in Section 5, while Section 6 shows the experimental results. The supplementary material accompanying this article contains derivations of the inference procedure described here, along with additional qualitative results. Source code and data for reproducing the results described here is available from [10].

## 2 Related Work

Here we briefly review the skin colour modelling literature, focusing on adaptive models, which are more relevant to the scope of this paper. We also briefly review the more recent work on hair segmentation.

**Skin Modelling** There is a strong body of literature on general skin colour modelling, but less so for adaptation of a colour model to an individual person. Kakumanu *et al.* [11] provide a recent and comprehensive survey of the field. They provide overviews of colour spaces, parametric and non-parametric models for colour distributions, public databases, illumination invariance and adaptation. However, they make no mention of adapting a general

skin colour model to an individual person.

One of the strongest colour models for classifying pixels into skin and background without adaptation is the histogram-based model of Jones and Rehg [10]. They provide a large labelled database of skin and background images, known as the Compaq skin database, and estimate the probability of observing each of the  $256^3$  possible RGB triples in skin and non-skin pixels. They then use Bayes' rule and the likelihood ratio between the two models to compute the probability that a given pixel is skin or not, thus avoiding having to set a skin likelihood threshold. When comparing colour spaces, Phung *et al.* [11] conclude that  $256^3$  histogram bins should be used for the general skin model (assuming that there are enough training images) in any 3-d colour space, but *not* in a reduced colour space, such as a 2-d chroma-only space. Luminance is important, as also pointed out by Kakumanu *et al.* [12].

**Skin Adaptation** Adaptation to illumination changes has been studied for both video and single images. However, the underlying assumptions are often different from the scope of this article — the focus is more on large and possibly abrupt illumination changes with respect to direction, intensity and colour, *e.g.* [13].

Adaptation of a general skin colour model to an individual person has mainly been considered in the context of face tracking. In the large majority of cases (*e.g.*, [14, 15, 16, 17, 18, 19, 20]), colour adaptation is presented as first learning a general colour model in the form of a Gaussian, Gaussian mixture or histogram, and then re-estimating the model parameters using selected data from a new image. Surprisingly, the use of a proper Bayesian approach as proposed here has not been considered. A typical example of the approach described above is in [21]. A single 2-d Gaussian is used for a general colour distribution, defined in an intensity normalised red–green colour space. The sample mean and covariance of pixels classified as skin in a new frame are used to adapt the mean and covariance of the Gaussian skin model using auto-regressive temporal smoothing.

The decision about which pixels to use for updating the skin model relies on two types of constraints, *viz.* chromatic and geometric. Soriano *et al.* [10] argue for the use of chromatic constraints rather than geometric constraints, and thus follow the approach of [21] summarised above, except that a skin histogram model is used. However, chromatic constraints can easily fail, either due to the lack of data for adaptation when the colour discrepancy with the general colour model is too great, or due to model drift in the presence of nearby skin-coloured hair, background or clothing regions. Geometric constraints have been used in a weak fashion for data selection — *e.g.*, by sampling the most reliable skin pixel locations given a face bounding box as provided by a tracker or a face detector [22, 23] — and to select background pixels when adapting a background colour model in order to apply a (desirable) colour likelihood ratio test rather than using a fixed threshold. However, in the latter case blond or red hair can often be confused with skin, since this reliability-driven geometric constraint will avoid selecting background pixels near the face.

In this work, we use both chromatic and geometric constraints in a principled Bayesian framework for colour distribution adaptation.

**Hair Colour Segmentation** More recently, a few papers have addressed the problem of hair segmentation [9, 24, 25, 26]. Most of these use similar approaches to the skin detection literature, but usually also exploit a segmentation algorithm, *e.g.* a graph-cut method in [24] and a snake-based active contour in [9]. Although they are computationally efficient, these methods do not have the same properties as our approach. Some use a simple colour model or geometric prior that allow them to work with a controlled background, or assume little variation in hair styles or pose, and they all lack a probabilistic framework for adaptation.




Figure 1: (a) Directed graph of the class and colour models. The coherent probabilistic index map (CPIM) prior is used to assign a discrete class label,  $z_{Ixy}$ , to each pixel in an image,  $I$ . One palette (that is, a distribution over colours) is drawn from a colour prior, for each class in each image. Each palette is a parametrised (by  $\Theta_{Ik}$ ) distribution — see Section 3.1 for more details. Given the class of a pixel and the palette of the class, the colour of the pixel,  $\tilde{c}_{Ixy}$ , is drawn from the corresponding palette. (b) Discrete colour histogram prior and colour likelihood. The palette for class  $k$  in image  $I$  is a normalised histogram with uniform bins, defined by a discrete probability vector,  $\tilde{\pi}_{Ik}^{(c)}$ , drawn from a Dirichlet prior,  $D(\tilde{\pi}_{Ik}^{(c)} | \tilde{\alpha}_k^{(c)})$ .

### 3 The Proposed Model

Fig. 1a shows an overview of the full generative model in the form of a directed graph. In the model, each image is generated as follows:

1. A palette is generated for each class. For each of the skin and hair classes, a normal distribution over colours in  $YP_bPr$  colour space is drawn from a normal–gamma prior. For each of the background and clothing classes, a discrete distribution over colour bins is drawn from a Dirichlet prior.
2. Each pixel is assigned a class from the coherent probabilistic index map (CPIM) prior, which is comprised of a PIM prior (*i.e.*, a discrete distribution over classes at each pixel) and a Markov random field between neighbouring pixels. This means that the classes of pixels are *not* drawn independently. See Section 3.2 for details.
3. For each pixel, a colour is generated from the palette corresponding to its class.

#### 3.1 Discrete and Continuous Palettes

Modelling a palette requires a choice of colour representation and a prior distribution over colours using the representation. Here we motivate our choices for different palettes.

**Discrete Palettes for Background and Clothing** For these classes we expect colour distributions to be broad and likely multi-modal since clothing and background are not constrained in terms of illumination or chroma. Therefore we use a discrete distribution (a normalised histogram) in RGB colour space. Formally, a discrete palette is a discrete distribution vector,  $\tilde{\pi}_{Ik}^{(c)}$ , over colour bins. The subscripts indicate that a palette belongs to class  $k$  in image  $I$ . The superscript  $(c)$  indicates that this discrete probability vector is for colour modelling. Fig. 1b shows a factor graph [1] of the discrete palette prior and colour likelihood. (Factor graphs explicitly show the variables, factors and factorisation properties of a probabilistic model. See Chapter 8 of [1] for a good introduction to factor graphs.) The probability that a colour in bin  $b$  is drawn is simply the  $b$ th entry in the vector,  $[\tilde{\pi}_{Ik}^{(c)}]_b$ , divided by the volume of the bin in colour space,  $v_{\text{bin}}$ , to normalise the distribution. This normalisation is important for consistency with the continuous palette model. The notation  $[\tilde{a}]_b \equiv a_b$  is used to index

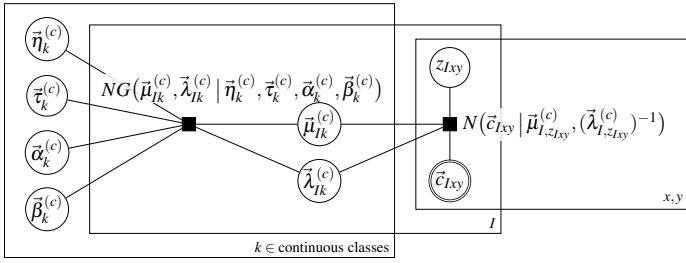


Figure 2: Continuous colour model and likelihood. The palette for class  $k$  in image  $I$  is a product of normal distributions over the three colour components — Y, P<sub>b</sub> and P<sub>r</sub> — parametrised by means,  $\vec{\mu}_{I_k}^{(c)}$ , and precisions (inverse variances),  $\vec{\lambda}_{I_k}^{(c)}$ . Each mean and inverse variance is drawn from a normal–gamma (NG) prior.

into a vector. The square brackets simply serve to aid legibility. The prior distribution over palettes, which can also be interpreted as the general distribution over colours for the class, is chosen to be the Dirichlet distribution, which is conjugate to the discrete distribution.

**Continuous Palettes for Skin and Hair** We model each of the skin and hair palettes using an axis-aligned normal distribution in YP<sub>b</sub>P<sub>r</sub> colour space. This space is used in the JPEG compression standard and was designed to separate lightness from colour. As reported in the literature, skin tone as represented by the chroma components (P<sub>b</sub> and P<sub>r</sub>) is very localised for skin in general [14] and we expect it to be even more localised for the skin of an individual person. Fig. 2 shows a factor graph of the continuous palette prior and colour likelihood. A continuous palette is a factored normal distribution over colour components, parametrised by a mean vector,  $\vec{\mu}_k^{(c)}$ , and a precision vector,  $\vec{\lambda}_k^{(c)}$ . We use the conjugate normal–gamma distribution as palette prior. The conjugacy makes for straightforward analytical inference and computationally efficient update equations — see Section 5 for details.

We would like to emphasise that, to the best of our knowledge, modelling general skin colour as a distribution over distributions has not been explored in the literature, where both the general and adapted skin colour models are usually handled as distributions over colours. The advantage of our approach is that estimating the adapted colour model is handled in a principled Bayesian framework by estimating the posterior belief over, for example, a person’s skin palette based on the prior (general) skin model and the colour likelihood function.

### 3.2 Coherent Probabilistic Index Maps (CPIM)

In this section, we first describe the probabilistic index map (PIM) of Jojic and Caspi [6], and then our extension using a Markov random field.

**Probabilistic Index Maps** The main objective of the PIM model is to learn a prior over the class of each pixel in an image. The idea is that the segmentation maps of (aligned) images of the same object are often consistent across instances, even if the pixels of the same class across training images are not characterised by the same colour distribution parameters — *i.e.*, they have different palettes. More formally, the PIM prior is defined as:

$$P_{\text{PIM}}(\langle z_{xy} \rangle | \langle \vec{\pi}_{xy}^{(z)} \rangle) = \prod_{x,y} [\vec{\pi}_{xy}^{(z)}]_{z_{xy}} \quad (1)$$

where  $\langle z_{xy} \rangle$  denotes the set of class label variables at all pixel coordinates, and  $\vec{\pi}_{xy}^{(z)}$  is a discrete probability vector over classes at pixel  $(x, y)$ . The PIM model assumes an independent




Figure 3: (a) The PIM model. The distribution over classes at each  $(x, y)$  coordinate is independent of the other coordinates. The class  $z_{I_{xy}}$  is drawn from a discrete latent distribution  $\bar{\pi}_{xy}^{(z)}$ . (b) CPIM: the coherent PIM model. The distribution over classes at each  $(x, y)$  coordinate depends on its four neighbours through a Markov random field.

prior over pixels, and hence spatial consistency is not incorporated. The prior parameters,  $\langle \bar{\pi}_{xy}^{(z)} \rangle$ , are estimated from images with ground truth class labels for each pixel.

**Incorporating a Markov Random Field** We know that adjacent pixels in an image belong to the same class with high probability. This knowledge is encoded in the prior over classes by incorporating into the PIM model a coherence energy function in the form of a Markov random field that encourages neighbouring pixels to be in the same class. The standard PIM prior (1) is modified as follows:

$$P_{\text{MRF}}(\langle z_{xy} \rangle | \langle \bar{\pi}_{xy}^{(z)} \rangle, k_{\text{MRF}}) \propto P_{\text{PIM}}(\langle z_{xy} \rangle | \langle \bar{\pi}_{xy}^{(z)} \rangle) e^{-E_{\text{MRF}}(\langle z_{xy} \rangle)} \quad (2)$$

where the Markov random field (MRF) energy function is

$$E_{\text{MRF}}(\langle z_{xy} \rangle) = -k_{\text{MRF}} \sum_{(xy, uv) \in \mathcal{C}} \mathbb{I}[z_{xy} = z_{uv}] \quad (3)$$

and  $\mathcal{C}$  denotes the set of pairwise cliques (*i.e.* the set of pairs of neighbouring coordinates  $(x, y)$  and  $(u, v)$  — we use a 4-connected neighbouring system) in the image and  $k_{\text{MRF}}$  is a constant indicating how important it is that adjacent pixels have the same class. This modified prior will be referred to as *coherent* PIM (CPIM). The factor graphs for PIM and coherent PIM are shown in Fig. 3.

## 4 Data

We used the Compaq skin database [2] for all of the training and evaluation in this article. The full database contains more than 13,000 images in two groups — skin and non-skin. Skin images have ground truth indicating which pixels contain skin, but this labelling was done conservatively in the sense that labels have very high precision (presumably 100%) but low recall. Therefore we can use the labels to construct a good skin colour model, but not to test the accuracy of skin segmentation. Non-skin images in the database were used to learn the background colour model prior, as described in the next section. The background prior was also used as the prior for clothing, since we expect that clothing can be any colour and since there is not enough labelled data to estimate a good prior for clothing. To learn the CPIM prior maps (and for testing), several images from the database containing faces detected by the Viola–Jones face algorithm [13] as implemented in the OpenCV library [6] were labelled for skin, hair, clothing and background. The size of a cropped image was chosen to be twice the size of the bounding box returned by the face detector, yielding a set of head-and-shoulder images, as illustrated in Fig. 4a.




Figure 4: (a) Example of an annotated image with skin (yellow), hair (red), clothing (green) and background (blue). White regions are unlabelled and usually close to boundaries where it is difficult to discriminate between classes. Unlabelled pixels were not used for training or testing. (b)–(e) The PIM prior that each pixel belongs to the skin, hair, clothing and background classes, respectively. Black indicates high probability and white, low probability.

## 5 Inference

Inference is done in two stages. Firstly, the prior parameters for CPIM and for each of the colour palettes are estimated from training data. Secondly, using these priors and given a test image, approximate Bayesian posterior beliefs over the class of each pixel and the colour distribution for each class are computed. The inference procedures for these stages are summarised below, with mathematical derivations, which are straightforward but lengthy, relegated to the supplementary material.

**PIM Prior** The PIM prior parameters,  $\langle \vec{\pi}_{xy}^{(z)} \rangle$ , are estimated directly from the fully labelled images by counting the number of times that each pixel belongs to each class in the training data. Left–right symmetry is enforced by flipping and averaging the estimated priors. Fig. 4b–e shows the resulting distribution over classes at each pixel in the image.

**Discrete Palette Prior** The Dirichlet prior for the discrete background palette (see Fig. 1b) is learned from all pixels in the non-skin images of the Compaq database. The discrete distribution for each image,  $\vec{\pi}_{jk}^{(c)}$ , can be integrated out analytically to get a Pólya distribution over the number of times that the colours in each histogram bin appear in each image. We estimate a regularised maximum likelihood setting of  $\vec{\alpha}_k^{(c)}$  for this distribution using the method of Minka [13]. The regularisation simply adds a small initial count to each histogram bin so as to avoid numerical singularities associated with zero counts. The mean of the Dirichlet prior over the background palette is visualised in Fig. 5a.

Due to a lack of labelled training data for clothing pixels, the prior distribution estimated for the background palette is also used as the prior for the clothing palette.

**Continuous Palette Prior** The prior parameters for continuous colour distributions are also estimated using a regularised maximum likelihood procedure. In this case the data likelihood is a bit more complex than for discrete distributions. First, we compute the maximum likelihood estimate for the normal distribution over colours in each image in the training data. Next, the maximum likelihood fit of the palette prior parameters is computed by numerical optimisation. The prior parameters ( $\vec{\eta}$ ,  $\vec{\tau}$ ,  $\vec{\alpha}$  and  $\vec{\beta}$ ) were introduced in Fig. 2. Each parameter vector has 3 entries for the 3 colour components. Since colour distributions are axis-aligned, we optimise one entry from each of these vectors (*i.e.*, 4 parameters) at a time.

The skin palette prior is estimated from the skin pixels in the face region of each Compaq image in which a face was detected. The hair palette prior is estimated from the much smaller set of fully labelled images from the same database. The marginals over colours from the




Figure 5: Marginal prior distribution over colours for the (a) background, (b) skin and (c) hair palettes. Each image represents a distribution over colours where the area devoted to the colour represents the probability mass of that colour. Colours are ordered in descending order of probability from top–left to bottom–right.

skin and hair prior palettes are shown in Fig. 5b and 5c, respectively.

**Adaptation and Segmentation in a Test Image** In a test image we compute approximate posterior beliefs over the class of each pixel,  $\langle z_{xy} \rangle$ , and the palette of each class using variational inference [8]. The main idea is to minimise the KL-divergence between the full joint posterior, which is difficult to compute, and a factorised belief over pixel classes and palettes, which is easy to compute. In our case the approximate posterior belief over the class of each pixel is another discrete distribution; the posterior belief over a discrete palette is a Dirichlet distribution; and the posterior belief over a continuous palette is a normal–gamma distribution. The approximate posterior is improved iteratively by updating one factor in the posterior belief at a time while keeping all other factors fixed. Using the implementation of [10] and with 20 variational iterations, running the inference procedure on a test image takes approximately 250ms.

## 6 Experiments

We conducted numerical experiments to validate the different elements of our model. The experimental setup is described below along with qualitative and quantitative results.

**Experimental Protocol** We compare methods based on accuracy (accu), *i.e.* the number of labelled pixels that were classified correctly, and also the precision (prec) and recall (rec) for each class, defined as:


$$\text{accu} = \frac{\sum_{c \in \text{class}} \text{tp}(c)}{\sum_{c \in \text{class}} [\text{tp}(c) + \text{fn}(c)]}, \quad \text{prec}(c) = \frac{\text{tp}(c)}{\text{tp}(c) + \text{fp}(c)}, \quad \text{rec}(c) = \frac{\text{tp}(c)}{\text{tp}(c) + \text{fn}(c)} \quad (4)$$

where tp, fp and fn are the true positive, false positive and false negative counts respectively. These measures are computed using all labelled pixels in the test set.

Since we have only 60 fully labelled images, experimental results are computed using cross-validation with 5 folds. In each fold, a fifth of the labelled data is held out as the test set and the other 4 fifths are used to learn the hair palette and PIM priors. Recall that the skin and background palette priors are learned using the large amount of skin and non-skin labelled data in the Compaq database, and that the clothing palette prior is set equal to the background palette prior. Discrete palettes have  $16^3 = 4096$  bins.

**Markov Random Field Results** We tested the importance of the Markov random field for improving the system. Turning off the MRF corresponds to setting  $k_{\text{MRF}} = 0$  and increasing





class	s (con)	h (con)	c (dis)	b (dis)
precision	0.951	0.820	0.725	0.853
recall	0.846	0.626	0.869	0.922
class	s (dis)	h (dis)	c (dis)	b (dis)
precision	0.919	0.780	0.746	0.876
recall	0.882	0.677	0.849	0.896

(b)

Figure 6: Markov random field experiments. (a) Overall accuracy as a function of the field strength,  $k_{\text{MRF}}$ . Blue circles show the results when using continuous skin and hair palettes, and red triangles when using discrete palettes. (b) Precision and recall values at  $k_{\text{MRF}} = 3$  for each class (skin, hair, clothing, background) under both conditions (continuous and discrete).

this constant corresponds to strengthening the MRF, meaning that neighbouring pixels are more likely to belong to the same class. Fig. 6 summarises the results of the experiment. The accuracy jumps by about 4% when the Markov random field is turned on ( $k_{\text{MRF}} = 1$ ), but then saturates quickly for larger values of the field strength. The maximum was found to be at  $k_{\text{MRF}} = 3$ . This shows that the MRF is an important part of the model.

**Continuous vs Discrete Palettes** To test the value of using continuous palettes, we also ran the above experiment with discrete palettes for the skin and hair classes, keeping the other experimental conditions the same as before. In this case the predictive accuracy is about the same, but the same dependence on the presence of the MRF is observed, as shown in Fig. 6a. The main difference lies in the precision and recall for each type of colour distribution. Note from Fig. 6b that recall for skin classification improves when using the discrete palette, but at the cost of precision. It appears that since the discrete palette is a more flexible representation than the continuous one, it adapts to cover more skin-coloured pixels, but that same flexibility also causes it to cover non-skin pixels. One should prefer one model type over the other based on the application of the overall model.

## 7 Conclusion

This paper addressed the joint segmentation and colour adaptation task. Our contributions are: a Bayesian method for computing and adapting hyper-beliefs over skin palettes (*i.e.*, distributions over distributions rather than expected distributions), the mixing of discrete and continuous palettes, the combination of a Markov random field prior with the probabilistic index map model, and a variational inference procedure. Numerical results demonstrated the validity of the approach. Future work includes the use of this model for head pose estimation and tracking. One limitation of our approach is the single spatial prior, which could be improved by considering mixture models, *e.g.* of different hair styles [□].




Figure 7: Results of segmenting two challenging images into skin (yellow), hair (red), clothing (green), and background (blue) regions. Pixels for which the class posterior probability was below 95% were left white. (a) Original image. (b) Segmentation with full model, using continuous skin and hair palettes. (c) Segmentation without the MRF. (d) Segmentation using discrete palettes. In the first image both the background and the face are skin-coloured, but the model is able to separate them well. In the second image the person has no hair and the shadow in the background forms a region that can be interpreted as hair. This fools the method into fitting the hair palette to this region.

## References

- [1] Adaptive face colour modelling and segmentation project page. URL <http://www.idiap.ch/software/facecolormodel/>.
- [2] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [3] D. Comaniciu and P. Meer. Mean shift: a robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619, 2002.
- [4] B. J. Frey. *Graphical Models for Machine Learning and Digital Communication*. MIT Press, 1998.
- [5] Intel Corporation. OpenCV, version 2.0.0, October 2009. URL <http://opencv.willowgarage.com/wiki/>.
- [6] N. Jojic and Y. Caspi. Capturing image structure with probabilistic index maps. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2004.
- [7] M. J. Jones and J. M. Rehg. Statistical color models with application to skin detection. *International Journal of Computer Vision*, 46:81–96, 2002.
- [8] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37:183–233, 1999.
- [9] P. Julian, C. Dehais, F. Lauze, V. Charvillat, A. Bartoli, and A. Choukroun. Automatic hair detection in the wild. In *International Conference on Pattern Recognition*, 2010.
- [10] P. Kakumanu, S. Makrogiannis, and N. Bourbakis. A survey of skin-color modeling and detection methods. *Pattern Recognition*, 40(3):1106–1122, 2007.

- [11] K. Lee, D. Anguelov, B. Sumengen, and S. B. Gokturk. Markov random field models for hair and face segmentation. In *Automatic Face and Gesture Recognition*, 2008.
- [12] C. Liensberger, J. Stottinger, and M. Kampel. Color-based and context-aware skin detection for online video annotation. In *IEEE International Workshop on Multimedia Signal Processing*, 2009.
- [13] T. Minka. Estimating a Dirichlet distribution. Technical report, 2009.
- [14] P. Perez, C. Hue, J. Vermaak, and M. Gangnet. Color based probabilistic tracking. In *European Conference on Computer Vision*, 2002.
- [15] S. L. Phung, A. Bouzerdoum, and D. Chai. Skin segmentation using color pixel classification: Analysis and comparison. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(1):148–154, 2005.
- [16] L. Sigal, S. Sclaroff, and V. Athitsos. Skin color-based video segmentation under time-varying illumination. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(7):862–877, 2004.
- [17] M. Soriano, B. Martinkauppi, S. Huovinen, and M. Laaksonen. Adaptive skin color modeling using the skin locus for selecting training pixels. *Pattern Recognition*, 36: 681–690, 2003.
- [18] P. Viola and M. Jones. Robust real-time face detection. *International Journal of Computer Vision*, 57(2):137–154, 2004.
- [19] D. Wang, X. Chai, H. Zhang, H. Chang, W. Zeng, and S. Shan. A novel coarse-to-fine hair segmentation method. In *Automatic Face and Gesture Recognition*, 2011.
- [20] Y. Yacoob and L. S. Davis. Detection and analysis of hair. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(7), 2006.
- [21] J. Yang and A. Waibel. A real-time face tracker. In *IEEE Workshop on Applications of Computer Vision*, pages 142–147, 1996.
- [22] T. Yoo and I. Oh. A fast algorithm for tracking human faces based on chromatic histograms. *Pattern Recognition Letters*, 20:967–978, 1999.