EUROGRAPHICS 2011 / M. Chen and O. Deussen
(Guest Editors)

Volume 30 (2011), Number 2

Computer-Suggested Facial Makeup

Kristina Scherbaum[1] Tobias Ritschel[2,3] Matthias Hullin[1] Thorsten Thormählen[1,5] Volker Blanz[4] Hans-Peter Seidel[1]

[1]MPI Informatik [2]Télécom ParisTech / LTCI / CNRS [3]Intel Visual Computing Institute [4]University of Siegen [5]University of Tübingen

Abstract

Finding the best makeup for a given human face is an art in its own right. Experienced makeup artists train for years to be skilled enough to propose a best-fit makeup for an individual. In this work we propose a system that automates this task. We acquired the appearance of 56 human faces, both without and with professional makeup. To this end, we use a controlled-light setup, which allows to capture detailed facial appearance information, such as diffuse reflectance, normals, subsurface-scattering, specularity, or glossiness. A 3D morphable face model is used to obtain 3D positional information and to register all faces into a common parameterization. We then define makeup to be the change of facial appearance and use the acquired database to find a mapping from the space of human facial appearance to makeup. Our main application is to use this mapping to suggest the best-fit makeup for novel faces that are not in the database. Further applications are makeup transfer, automatic rating of makeup, makeup-training, or makeup-exaggeration. As our makeup representation captures a change in reflectance and scattering, it allows us to synthesize faces with makeup in novel 3D views and novel lighting with high realism. The effectiveness of our approach is further validated in a user-study.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture I.5.4 [Computer Graphics]: Applications—Computer vision

1. Introduction

For many centuries people have changed facial reflectance by different means to achieve pleasant facial appearance. In current everyday life, facial makeup is mostly used to mimic healthy and attractive facial geometry and reflectance. Finding the best makeup for a given human face can be considered an art form, and nowadays many people trust the advice of professional makeup-artists.

This paper introduces a model of facial makeup using a database of example faces with and without makeup. The makeup model allows, for example, to provide computer-suggested makeup for new subjects, which are not part of the database. It is also possible to rate an applied makeup in order to give an objective feedback. For all these applications we make the assumption that the choice of human facial makeup is based on the more or less conscious process of mapping facial features to reflectance and scattering changes. This is of course just an approximation of reality, as there are some widely accepted standards but no strict rules for makeup [Tho05]. Furthermore, the mapping from facial appearance to makeup might change over culture, ethnics, and history. To be practical, we limit our investigations to western, 21st century, female facial makeup.

In contrast to existing approaches [GS09, TTBX07], our approach does not only allow to transfer makeups between subjects but provides a computer-suggested makeup. Furthermore, we perform our makeup analysis and synthesis in 3D, whereas related approaches worked in the 2D domain. Nevertheless, our 3D analysis can provide makeup suggestions, if only a 2D photograph of the subject is available. The proposed 3D synthesis allows to simulate the suggested makeup under different lighting conditions.

The paper is structured as follows: After reviewing previous work in Section 2, we describe our approach to build a makeup model in Section 3 that leads to a number of example applications described in 4 that are validated in a perceptual study. After a discussion 5, we conclude in Section 6.

2. Related Work

Facial Appearance Convincingly modeling the appearance of the human body and especially the human face has been a long-standing goal in computer graphics. Modeling facial appearance, which consists both of geometry as well as reflectance and scattering, is challenging, because human observers are well-tuned to perceive faces [HHG00] and even subtle imperfections may easily lead to an unpleasant impression [SN07]. To capture the essence of facial geometry, Blanz and Vetter [BV99] introduced a morphable face model. Besides geometry, the appearance (i.e. reflectance and scattering) of human skin was captured [MWL*99, WMP*06]. The effect of aging, alcohol consumption and also foundation cosmetics can be predicted by a model based on the separation of hemoglobin and melanin [TOS*03].

Makeup The main purpose of makeup is to temporarily change the facial appearance. Most people apply makeup to either cover unwanted facial structures, such as wrinkles or pores, or to emphasize specific features. Others apply makeup to express themselves and to play around with different styles and looks. However, makeup is a very subjective matter. The common understanding of makeup varies between different cultures and changed over history [Cor72]. Further, cosmetic makeup is, and has been historically, much more common among women [Cor72] and we will hence limit our investigation to female, present time makeup for simplicity. Russel [Rus03] studies the relation of contrast between different facial parts, indicating that luminance changes as found in cosmetics are more effective on female faces.

Given a certain facial appearance, makeup is used to emphasize certain features or deemphasize others. Therefore, makeup is a change of facial appearance that depends on the initial appearance. This change is not limited to color, but also gloss and scattering are affected. To our knowledge, makeup was neither rendered nor captured in previous work and only considered in an image-based context [TTBX07, GS09]. In this work, we will argue how makeup is a change of facial appearance as a function of facial appearance, i. e. it is a mapping, that can be modeled given some examples.

Adding or emphasizing existing makeup can approximately be achieved for live footage using simple image filters [NRK99]. Both Numata et al. [NRK99] and Tsumura et al. [TOS*03] change the melanin and hemoglobin mixture to indirectly mimic the effect of foundation cosmetics, which in reality does not change the melanin and hemoglobin mixture.

The system of Tong et al. [TTBX07], allows to transfer cosmetic when given a before-after image pair. Guo and Sim [GS09] transfer makeup and other facial details from a single image to another image but without an explicit notion of makeup itself. Finding a suitable "decoration" of a face is similar to automated caricature as done by Chen and colleagues [CXS*01].

Facial Attractiveness Facial attractiveness is perceived consistently between observers of different age, sex and cultural background [LR90] and is believed to relate to averageness [LR90], symmetry [PBP*99] and sexual dimorphism [PLP*98].

Mappings *from* human shape to simple attributes were investigated by [EDR06] (attractiveness) or [WV09] (social judgments). The model of Eisenthal et al. [EDR06] was later used for facial image beautification by Leyvand and coworkers [LCODL08]. While all previous work has analyzed real attractiveness (and synthesized virtual beauty in the case of Leyvand's [LCODL08] work), we conduct an analysis, complemented by a synthesis step that improves real physical attractiveness when applying our suggested makeup.

3. Analysis and Synthesis of Makeup

To extract the essence of makeup, we model it as a mapping from a domain of facial appearance to a range of appearance changes, i.e. makeup. To this end, we first acquire facial appearance including makeup then, extract the makeup, construct the mapping, and finally generate a suggested makeup.

3.1. Facial Appearance Capture

We acquired a database of $N = 56$ female faces in two states: without makeup and with professional makeup. Subjects were placed in a light tent surrounded by six projectors. To capture advanced facial properties, we project a number of structured light patterns in each state: Gradients along X, Y and Z and vertical stripe patterns of increasing frequencies (8 octaves) (cf. Fig. 1). The patterns were successively projected in less than a second and each was captured using a *Canon EOS 5D Mark II* camera. First, we used the stripe patterns to perform a coarse point-based 3D reconstruction of the subject's face surface. Next, we fit a morphable face model [BV99] both to all "full-on" (cf. Fig. 1) images with and without makeup. Using the registered face model we have converted all images into a common parameterization, i. e. we make sure that the tip of the nose is in the same location in all images (Fig. 2). Finally, we use manual 2D image deformation to make the registration pixel-accurate. Remaining holes,

Figure 1: *The controlled-light setup used for facial appearance capture.*

Figure 2: *Registered diffuse appearance and makeup.*

i. e. parts of the face not seen in the image are filled using texture in-painting [Tel04]. Please note that these manual steps are performed while building the model, but no manual intervention is required when using the model.

From the gradient images we acquire photometric normals [MHP*07]. We did not use polarized light to separate specular and diffuse normals. Finally, we remove low-frequency normal bias using Nehab et al.'s [NRDR05] method.

We use the 3D surface model to perform inverse lighting i.e. using the "full-on" image (cp. Fig. 1) to compute spatially varying diffuse reflectance [WMP*06]. An individual specularity and glossiness term for each subject is computed from the "single light source" image and for different segments, such as the lips. The change of specularity and glossiness after applying makeup is most noticeable on the lips. Fig. 3 shows a typical example, where the lips show stronger specularity and gloss after the makeup is applied.

Finally, we compute spatially varying subsurface scattering strength from the stripe patterns by direct-indirect separation [NKGR06]. We approximate scattering using a sum of Gaussians, which has shown to be perceptually plausible [JSG09]. Comparing the subsurface scattering with and without makeup, it can be observed that makeup reduces subsurface scattering of the skin. The makeup on the lips (typically lipstick) mostly eliminates the subsurface scattering.

In summary, for subject i in state X we acquired an *appearance image* $A_{i,X}$ that stores per-pixel: diffuse color RGB (3), 3D position XYZ (3), 3D normal XYZ (3), specularity (1),

Figure 3: *The measured specularity and glossiness changes after applying makeup. This is most noticeable on the lips: (Left) rendered lips without makeup, (Right) rendered lips with makeup.*

Figure 4: *Appearance images (here: a single subject) have a common parametrization for all subjects and all bands.*

glossiness (1), scattering strength RGB (3), and scattering size RGB (3) in a common (for all faces) parameterization (cp. Fig. 4). The values in brackets state the number of required bands. Thus, an appearance image has a total of 17 bands.

3.2. Makeup

Once acquired, we compute the change in facial appearance and call this change the *makeup*. The makeup of subject i from state X (without makeup) to state Y (with makeup) is denoted as $M_{i,X\to Y}$. Further, we define the ratio of appearance with and without makeup to produce makeup $M_{i,X\to Y} = \frac{A_{i,Y}}{A_{i,X}}$. Consequently, the multiplication of appearance without makeup $A_{i,X}$ and makeup $M_{i,X\to Y}$ produces appearance with makeup $A_{i,Y} = A_{i,X} \cdot M_{i,X\to Y}$. Note, that M_i and A_i are images and multiplication or division must be done per pixel and band. We assume that makeup has no geometric effects, i.e. normals and positions are not affected by makeup and are omitted. The change in diffuse color is most important and we use the ratio in RGB space [LSZ01] to express the effect. Other color spaces like LAB performed worse. Glossiness and specularity is modeled as scalar addition, all changes to scattering as monochromatic multiplication.

3.3. Appearance-to-Makeup Mapping

Let $\mathcal{A}$ be the space of all possible facial appearances and $\mathcal{M}$ the space of all possible makeups. We want to find the best mapping from a facial appearance $A \in \mathcal{A}$ to a makeup $M \in \mathcal{M}$. For each of our 56 examples in the database this mapping is given, because a particular makeup M_i to go from a no-makeup-state X into a makeup-state Y is only a function of the subject's appearance A_i. If we want to determine a makeup M_{query} for a new subject A_{query}, which is not in our database, we can perform nearest neighbor matching in the space of facial appearance, i. e.

$$M_{\text{query}} = M_j \tag{1}$$

where

$$j = \operatorname*{argmin}_i d(A_{\text{query}}, A_i). \tag{2}$$

A naïve distance function $d(A_{\text{query}}, A_i)$ for nearest neighbor matching would be the sum of absolute differences of all pixels and bands of the two appearance images. However, this approach would assume that each pixel and each band contains the same amount of information. Instead, we perform a PCA of all facial appearances A_i using the classical Eigenface approach [TP91]. The only difference is that compared to the classical approach our appearance images contain a larger number of bands, as for example the 3D coordinates. The Eigenface approach allows to use the Mahalanobis distance of the corresponding facial appearance PCA coefficients as the distance function $d(A_{\text{query}}, A_i)$. By determining the nearest neighbor in PCA space, it is ensured that the most descriptive pieces of information are used to differentiate between the subjects and the subject with the most similar appearance is selected (for queries within the database: "leave one out").

However, we found during our experiments that the nearest neighbor matching using the Mahalanobis distance of the PCA coefficients does only to some extent find the makeup that is most aesthetic for a given face. A human observer focuses strongly on certain features of the face, in particular on the color of the eyes, hair, or skin. These inter-subject differences can be also found in the PCA coefficients, but are not weighted as high in the Mahalanobis distance as they are weighted by a human observer. Consequently, we extend our distance function for nearest neighbor matching with the following heuristic:

$$d(A_{\text{query}}, A_i) = w_1\, d_{\text{pca}} + w_2\, d_{\text{eye}} + w_3\, d_{\text{skin}} + w_4\, d_{\text{hair}} \tag{3}$$

where d_{pca} is the Mahalanobis distance between the PCA coefficients of A_{query} and A_i, and d_{eye}, d_{skin}, and d_{hair} are the distances between the eye, skin, and hair color vectors of A_{query} and A_i in RGB space, respectively. The eye, skin, and hair color for each subject can easily be extracted from the rectified appearance images. We set the weights to $w_1 = 0.6$, $w_2 = 0.2$, $w_3 = 0.1$, $w_4 = 0.1$ in all our experiments. Experimenting with several combinations we found these heuristic weight parameters to yield the most convincing results.

Figure 5: *The two images on the right side show the query face after makeup transfer of the best match. By applying PCA compression our approach allows the re-synthesis of makeup that automatically omits unwanted transfer of personal details, like freckles or moles. (The makeup is applied 5 times for better visibility.)*

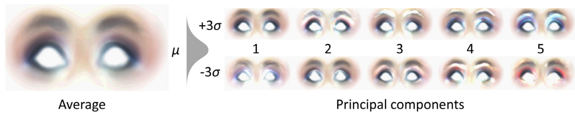

Figure 6: *The first five Eigen-makeups for the eye segment and the respective average. (The contrast was enhanced for better visibility.)*

3.4. Re-Synthesis of Makeup

Now that we have determined by nearest neighbor matching which makeup we want to copy, we could apply this makeup to the query appearance with $A_{\text{query},Y} = A_{\text{query},X} \cdot M_{\text{query},X\to Y}$. However, as shown in Fig. 5 this would result in a transferal of personal details, like freckles or moles. If the nearest neighbor subject has freckles or moles and these are covered by makeup, these details appear as *inverted* freckles or moles in the transfered makeup of the query subject. Thus, we need to find a way to regularize our makeup. In other words, we want to transfer only the essence of that particular makeup compared to other makeups, but not the personal details. This can be achieved by applying a PCA on all makeups M_i. Again, we use the Eigenface approach and calculate the corresponding *Eigen-makeups* (cp. Fig. 6). The resulting $N = 56$ Eigen-makeups are sorted according to their eigenvalues. The Eigen-makeups with large eigenvalues are likely to contain variations that can be observed between many makeups, whereas the Eigen-makeups with small eigenvalues contain the personal details. We then re-synthesis each makeup $\tilde{M}_i$ by using the 10 Eigen-makeups with the 10 largest Eigenvalues. As shown in Fig. 5 these makeups are regularized to omit personal details and can be applied to the the query appearance with

$$A_{\text{query},Y} = A_{\text{query},X} \cdot \tilde{M}_{\text{query},X\to Y} \tag{4}$$

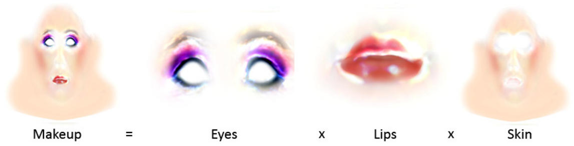

Figure 7: *We compose the makeup merging three regions (eyes, lips, skin). Each of the regions is reconstructed from its own Eigenspace using 35,15,10 Eigenvectors respectively. Note that the shown makeup is squared for better visibility.*

We can improve upon this approach by performing individual PCAs for different parts of the face. This allows the PCA to identify those differences between makeups that are typical for these different parts. We segment the face into eyes, lips as well as skin and create individual Eigenspaces for each makeup segment (cf. Fig. 7). The eye and lip regions have soft borders to fade into white, i.e. the identity makeup. We then re-synthesize a complete makeup $\tilde{M}_i$ by re-synthesizing each part individually and compose the individual regions to a complete makeup. As a result, the regularization is performed on the individual parts rather than the whole makeup, which we found to perform better in practice. Using the new approach we can for example automatically remove wrinkles all over the face while preserving the full variance of possible eye makeups.

3.5. Rendering

Having acquired the facial appearance from the subject in question, we can show a visualization of novel synthesized makeup from novel viewpoints under novel lighting, including effects such as subsurface scattering [JSG09] and specular lighting in real-time.

4. Applications and Results

4.1. Applications

Computer-Suggested Makeup The key application of our model is computer-suggested makeup. Given a 2D image (e.g., a photograph) of a subject as a query, we fit the morphable face model [BV99] to the face in the image. This gives us the facial appearance without makeup $A_{X,\text{query}}$. However, this appearance image has only 6 bands (diffuse color RGB and position XYZ) in contrast to the 17 bands we have acquired for the appearance images A_i in the database, which were captured in the light tent. Consequently, we can use only the first 6 bands to perform nearest neighbor matching, as described in Section 3.3 (cp. Fig. 8). Once the nearest neighbor is found, we just copy the missing bands from the nearest neighbor to the query appearance, except the normals. This allows us to perform convincing real-time 3D renderings of the subject without makeup under varying view position and lighting setups. We then re-synthesize the suggested makeup from the nearest neighbor makeup (cp. Section 3.4) and generate $A_{Y,\text{query}}$. The user can now inspect the query subject

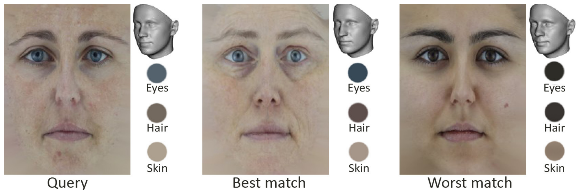

Figure 8: *The computer-suggested makeup is found by a nearest neighbor search that involves the 3D-shape PCA coefficients of the facial appearance A_i as well as skin-, eye-, and hair-color.*

with suggested makeup in real-time under varying viewpoints and different illuminations (Fig. 9).

Makeup Transfer Makeup transfer (as done image-based by Guo and Sim [GS09]) can be done more robustly in 3D. The 3D information allows to perform robust inverse lighting and, thus we can render the subject under different viewpoints and with novel lighting The disadvantage of our approach is that we require a with-and-without-makeup image pair to capture the makeup. The advantage of our approach is, that we use the PCA compression as a regularizer, whereas Guo and Sim's approach relies on low-pass filters to remove personal details that should not be transferred.

Our PCA approach allows to transfer finer details that are shared by all subjects, i. e. in proximity of the eyes, which are blurred away even for advanced filters (bilateral, non-local means).
To perform a makeup transfer, the user provides a facial appearance of one subject without makeup $A_{\text{source},X}$, the same subject with makeup $A_{\text{source},Y}$, and a second subject without makeup $A_{\text{target},X}$. As we already know which makeup we need to transfer, we can skip the nearest neighbor search and just need to perform the re-synthesis step from Section 3.4 to generate $A_{\text{target},Y}$. A user can then inspect the effect of the first subject's makeup on the second subject.

Automatic Rating of Makeup A variant of computer-suggested makeup is the automatic rating of existing makeup. Here, a user provides a with-and-without-makeup appearance pair $(A_{X,\text{query}}, A_{Y,\text{query}})$ to the system. Given this appearance pair the query makeup is given by

$$M_{\text{query}} = \frac{A_{\text{query},Y}}{A_{\text{query},X}} \tag{5}$$

In contrast to the computer-suggested makeup the nearest neighbor search is now performed in the makeup space $\mathcal{M}$ and not in the appearance space $\mathcal{A}$. To be more specific, we use the Mahalanobis distance of the makeup PCA coefficients as the distance function $d(M_{\text{query}}, M_i)$ to find the nearest neighbor in makeup space (instead of $d(A_{\text{query}}, A_i)$ in appearance space as before). Once we found the nearest neighbor makeup M_{nearest}, we know the corresponding appearance$A(M_{\text{nearest}})$ in the database. We can then return the Mahalanobis distance $d(A(M_{\text{nearest}}), A_{X,\text{query}})$ in appearance

Figure 9: *Our main application is computer-suggested makeup. Given a subject (One in each colum), starting from a face without makeup (First Row), we suggest a makeup (Second row), according to a mapping acquired from a makeup artist (his result shown in the third row). We include the makeup that fits the least (Fourth row) for comparison.*

space as a rating for the makeup. Thereby, a small distance is considered a better makeup.

Inverse Computer-Suggested Makeup In inverse computer-suggested makeup, our system generates a facial appearance A that would be best for a given makeup M_{query}. The approach is similar to the automatic rating of makeup, except that $A(M_{\text{nearest}})$ is returned by the system instead of a rating. This can be used for didactic purposes to study what facial appearance corresponds to which makeup (cp. Fig. 10).

Makeup Exaggeration Given a facial appearance without makeup A and its makeup M our model can be used for (de)-exaggeration. The naïve approach to this problem is to multiply M by a constant, which indeed increases the effect of M. However, best results are achieved when moving away form the intra-subject average (Fig. 11).

Figure 10: *Given a query Makeup (Left), the user can now search for the facial appearance (Right) that would fit best to his makeup. This could be useful training makeup artists.*

4.2. Perceptual study

We evaluated the effectiveness of our model for the "computer-suggested makeup" application in a perceptual

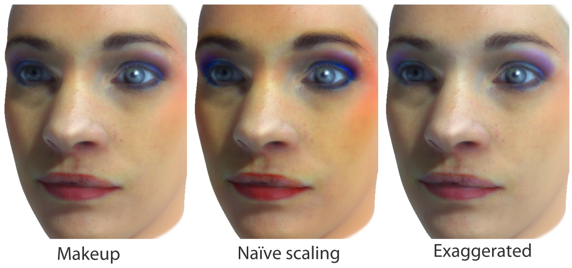

Figure 11: *Starting from a facial appearance with makeup (Left), naïve makeup scaling (Middle) leads to an overall scaled makeup, e.g. eyeshadow – which are found in many makeups – are enhanced. When scaling the Eigenmakeup instead ($\sigma = 2.0$), features typical for this individual makeup are exaggerated (Right).*

study. Participants compared three screenshots of rendered faces that were placed next to each other in one row: I_q (query subject without makeup), I_a (this subject with makeup), and I_b (this subject with a different makeup). The query image of the subject without makeup was placed in the middle of the both images with makeup.

We asked the same questionat all times: "What makeup fits best for the query appearance in I_q: The makeup in image I_a or the makeup in image I_b?". In each trial we randomized the position of the query image I_a and I_b. We ran three experiments, which used

1. the re-synthesized makeup applied by the professional makeup artist for I_a and a randomly selected makeup for I_b,
2. our suggested makeup for I_a and the farthest neighbor makeup for I_b. (The farthest neighbor makeup is the makeup with the largest distance $d(A_{query}, A_i)$),
3. our suggested makeup for I_a and a randomly selected makeup for I_b.

For each of these three experiments the participants were shown 17 image triplets. Thus, in total, each participant had to rate 51 image triplets.
As we ran the three experiments on the Amazon Mechanical Turk platform, we decided for two mechanisms avoiding random answers. First, within each set of 17 images we showed three images twice. Consequently each user rating these three images differently was excluded from the final results. Secondly, we evaluated the time users needed to answer the three experiment tasks. Heuristics showed 30 seconds to be the minimal time for answering a single task. Thus, users taking less than 30 seconds were excluded as well. Overall these restrictions reduced the number of valuable results from initially 210 to 146, which is 69%. Another side effect of the filtering is that the distribution of male and female participants differs in each experiment.

Table 1 summarizes the results of the three experiments. It can be observed that 62 % of all participants preferred the professional makeup over a random makeup. If we formulate the null hypothesis that professional and random makeup are equal, this null hypothesis can be rejected with a Pearson's chi-square test. The probability of the observed rating under the null hypothesis is $< 10^{-8}$ (p-value). Consequently, it is extremely likely that professional and random makeup are not equal. In the second experiment, 67 % of all participants preferred our suggested nearest neighbor makeup over the farthest neighbor makeup. Again, it can be shown with Pearson's chi-square test that this difference is statistically significant (p-value $< 10^{-15}$). Finally, in the third experiment, 58 % of the participants liked our suggested makeup better than a randomly selected makeup (p-value $< 10^{-4}$).

	Experiment 1		Experiment 2		Experiment 3	
	prof.	rand.	NN	FN	NN	rand.
All	62 %	38 %	67 %	33 %	58 %	42 %
Female	64 %	36 %	66 %	34 %	62 %	38 %
Male	58 %	42 %	67 %	33 %	55 %	45 %
Signif.	$p < 10^{-8}$		$p < 10^{-15}$		$p < 10^{-4}$	

Table 1: *Results for the three experiments of the study are shown. The upper 3 rows state the percentage of participants that voted for a particular makeup (prof. = professional makeup, rand. = random makeup, NN = nearest neighbor makeup, FN = farthest neighbor makeup). The last row shows the p-values of the corresponding Pearson's chi-square test.*

These results allow some interesting implications. The small advantage a professional makeup expert can improve upon randomness is at best $62\% - 50\% = 12\%$ and serves as a baseline of what can be achieved. Overall, the computer-suggested makeup does perform only slightly worse than the professional makeup (58 % vs. 62 %). This can be considered a very good result. Much higher percentages cannot be expected as we have modeled our mapping from the professional makeup which constitutes an upper bound. Examples for the stimuli used as well as the full ratings are found in the supplemental material.

5. Discussion

In all results presented, we use nearest-neighbor sampling over $\mathcal{M}$, though more advanced reconstructions can improve the results. Further, we use multiplication [LSZ01] of diffuse colors to simulate the effect of makeup, while the more advanced Kubelka-Munk-theory could be used. Linear blending of multiple components using the first three Eigenvectors that are sufficiently orthogonal creates in-between makeups, but becomes implausible for others. Here, more than 56 samples would be required. Assuming uniform distribution, we can infer from the given Eigenvector distribution that our makeup space is sufficiently dense for lips and skin while for the eyes more samples would be needed. Within our current setup we expect 100 samples to be suitable. In our approach, we currently assume that there is no bad example in the database and the expert was always right. To generalize to arbitrary non-expert input (e. g. from community photo databases), a rating of the individual makeup would need to complement the approach. We base our model on makeup performed by a single professional expert makeup artist. This introduces the artist's personal taste as a bias which could be reduced by using more artists. Consequently this would require even more samples in order to guarantee a densely sampled makeup space. We re-synthesize new makeups always from a single best-match makeup. This is on the one hand suitable as makeup artists focus on creating styles, where eye, skin and lip makeup homogeneously match. However, combining different makeups could lead to new and interesting makeup styles, such as mixing the facial regions of different makeups. In this case more sophisticated blending algorithms would be mandatory.

6. Conclusion and Future Work

We have presented a first step to computationally model the relation of facial appearance and facial makeup. Using a database of examples, we built a model that allows applications such as machine-suggested, individualized makeup that can improve perceived attractiveness as shown in a perceptual study. Despite previous approaches our machine-suggested makeup involves 3D information for both, analysis and synthesis. Such, results can be relighted and inspected in arbitrary poses and under arbitrary lighting conditions.

In future work, the ideas presented in this paper could be extended to model many other aspects of human "style" that might depend in a more or less observable way on individual human properties. Examples are hairstyle from facial appearance and color or clothing style from body appearance or human gait. Also external parameters could be included: makeup for a sunny day might have other requirements than the one for a nightclub.

References

[BV99] Blanz V., Vetter T.: A morphable model for the synthesis of 3D faces. In *Proc. SIGGRAPH* (1999), pp. 187–194. 2, 5

[Cor72] Corson R.: *Fashion in makeup: From Ancient to Modern Times*. Peter Owen Publishers, 1972. 2

[CXS*01] Chen H., Xu Y.-Q., Shum H.-Y., Zhu S.-C., Zheng N.-N.: Example-based facial sketch generation with non-parametric sampling. *Proc. ICCV 2* (2001), 433. 2

[EDR06] Eisenthal Y., Dror G., Ruppin E.: Facial attractiveness: Beauty and the machine. *Neural Comput. 18*, 1 (2006), 119–142. 2

[GS09] Guo D., Sim T.: Digital face makeup by example. In *Proc. CVPR* (2009). 1, 2, 5

[HHG00] Haxby J. V., Hoffman E., Gobbini M. I.: The distributed human neural system for face perception. *Trends Cogn Sci. 4* (2000), 223–233. 2

[JSG09] Jimenez J., Sundstedt V., Gutierrez D.: Screen-space perceptual rendering of human skin. *ACM Trans. on Applied Perception* 6, 4 (2009). 3, 5

[LCODL08] Leyvand T., Cohen-Or D., Dror G., Lischinski D.: Data-driven enhancement of facial attractiveness. *ACM Trans. Graph. (Proc. SIGGRAPH) 27*, 3 (2008). 2

[LR90] Langlois J. H., Roggman L. A.: Attractive faces are only average. *Psychological Science 1* (1990), 115–121. 2

[LSZ01] Liu Z., Shan Y., Zhang Z.: Expressive expression mapping with ratio images. In *Proc. SIGGRAPH* (2001), pp. 271–276. 4, 7

[MHP*07] Ma W.-C., Hawkins T., Peers P., Chabert C.-F., Weiss M., Debevec P.: Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proc. EGSR* (2007), pp. 183–194. 3

[MWL*99] Marschner S. R., Westin S. H., Lafortune E. P. F., Torrance K. E., Greenberg D. P.: Image-based BRDF measurement including human skin. In *Proc. EGWR* (1999), pp. 139–152. 2

[NKGR06] Nayar S. K., Krishnan G., Grossberg M. D., Raskar R.: Fast separation of direct and global components of a scene using high frequency illumination. *ACM Trans. Graph. 25*, 3 (2006), 935–944. 3

[NRDR05] Nehab D., Rusinkiewicz S., Davis J., Ramamoorthi R.: Efficiently combining positions and normals for precise 3D geometry. *ACM Trans. Graph. (Proc. SIGGRAPH) 24*, 3 (2005). 3

[NRK99] Numata K., Ri K., Kira K.: HD E-make: A real-time HD skin-make-up machine. *Eizo Joho Media Gakkai Gijutsu Hokoku 24* (1999), 45–50. 2

[PBP*99] Perrett D. I., Burt D. M., Penton I., Lee K., Rowland D., Edwards R.: Symmetry and human facial attractiveness. *Evolution & Human Behavior* (1999), 295–307. 2

[PLP*98] Perrett D., Lee K., Penton I., Rowland D., Yoshikawa S., Burt D. M., Henzi S. P., Castles D., Akamatsu S.: Effects of sexual dimorphism on facial attractiveness. *Nature* (1998), 884–887. 2

[Rus03] Russel R.: Sex, beauty, and the relative luminance of facial features. *Perception 32* (2003), 1093–1107. 2

[SN07] Seyama J., Nagayama R. S.: The uncanny valley: Effect of realism on the impression of artificial human faces. *Presence: Teleoper. Virtual Environ. 16*, 4 (2007), 337–351. 2

[Tel04] Telea A.: An image inpainting technique based on the fast marching method. *J. Graphics Tools 9*, 1 (2004), 23–34. 3

[Tho05] Thomson K.: *Sometimes Less Is Not More: A Fabulous Guide to Fabulous Makeup*. Lipstick Press, 2005. 1

[TOS*03] Tsumura N., Ojima N., Sato K., Shiraishi M., Shimizu H., Nabeshima H., Akazaki S., Hori K., Miyake Y.: Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin. In *Proc. SIGGRAPH* (2003), pp. 770–779. 2

[TP91] Turk M., Pentland A.: Face recognition using eigenfaces. *Proc. CVPR* (1991), 586–591. 4

[TTBX07] Tong W.-S., Tang C.-K., Brown M. S., Xu Y.-Q.: Example-based cosmetic transfer. In *Proc. PG* (2007), pp. 211–218. 1, 2

[WMP*06] Weyrich T., Matusik W., Pfister H., Bickel B., Donner C., Tu C., Janet McAndless J. L., Ngan A., Jensen H. W., Gross M.: Analysis of human faces using a measurement-based skin reflectance model. *ACM Trans. Graph. (Proc. SIGGRAPH)* (2006), 1013–1024. 2, 3

[WV09] Walker M., Vetter T.: Portraits made to measure: Manipulating social judgments about individuals with a statistical face model. *J. Vision 9*, 11 (2009), 1–13. 2