

Inverse Rendering with a Morphable Model: A Multilinear Approach

Oswald Aldrian
oswald@cs.york.ac.uk

William A. P. Smith
wsmith@cs.york.ac.uk

Department of Computer Science
The University of York, UK

Abstract

Well known results in inverse rendering show that recovery of unconstrained illumination, texture and reflectance properties from a single image is ill-posed. On the other hand, in the domain of faces linear statistical models have been shown to efficiently characterise variations in face shape and texture. In this paper we show how the inverse rendering process can be constrained using a morphable model of face shape and texture. Starting with a shape estimate recovered using the statistical shape model, we show that the image formation process leads to a system of equations which is multilinear in the unknowns. We are able to estimate diffuse texture, specular reflectance properties, the illumination environment and camera properties from a single image. Our approach uses relaxed assumptions and offers improved performance in comparison to the current state-of-the-art morphable model fitting algorithms.

1 Introduction

The appearance of a face in an image is determined by several underlying factors. The individual contributions can be divided into two groups: intrinsic facial features (shape and surface reflectance properties) and extrinsic parameters such as camera properties, pose and illumination. An image is a projection of a 3D world to a 2D view-plane. The forward process is well understood using principles borrowed from projective-geometry and biophysics [5] and successfully applied in computer graphics.

The inverse process on the other hand is ill-posed, especially when having only pixel intensity values from single images at hand. The same appearance can be explained in multiple different ways. For example, a face which appears red can be caused by viewing a pale face under red illumination or a red face under white illumination. In the context of morphable model fitting, the majority of previous work [3, 10] has attempted to explicitly model the illumination environment as a combination of ambient and one or more directional sources. Where complex illumination has been considered, such as by Zhang and Samaras [13], it has required restrictive assumptions to be made, namely that the reflectance properties are Lambertian. Shim *et al.* [12] model specularities in 2D using an empirical PCA model. They use the first principal component to model specularities. Empirical models are trained to specific imaging conditions and lack the ability to generalise to images taken under novel conditions.

For an object of known shape under complex natural illumination, the well known work of Ramamoorthi [R] shows how the spherical harmonic domain can be used to estimate one or more of: illumination, surface texture and reflectance properties. We revisit this classical formulation in the context of morphable models. Here, the model provides a linear basis for surface texture. Moreover, the texture parameter is constrained by a probabilistic prior. To further regularise the problem, we design a number of further constraints on the illumination environment: a grayworld constraint (which prefers similar spherical harmonic illumination parameters across colour channels) and a classical regularisation constraint (which prefers simpler illumination environments).

Our approach makes the least restrictive assumptions of any existing method for fitting a morphable model. Illumination is allowed to vary arbitrarily, specular reflectance is constrained only to be homogenous and symmetric about the reflection direction and the gain, contrast and offset of the camera can be unknown. Despite this, our formulation is convex and therefore both efficient and guaranteed to obtain the globally optimal solution.

2 3D Morphable Model

A 3D morphable model is a statistical model of shape and texture. Given j face meshes which are in dense correspondence, each mesh is written as a vector $\mathbf{m} = [x_1 \ y_1 \ z_1 \ \dots \ x_p \ y_p \ z_p]^T \in \mathbb{R}^n$ where p is the number of vertices and $n = 3p$. Principal Components Analysis (PCA) is applied to the data, providing a mean face, $\bar{\mathbf{m}}$, and a set of $j - 1$ eigenvectors, the i th of which is denoted \mathbf{M}_i with corresponding variance $\sigma_{m,i}^2$. A linear surface texture model is constructed in exactly the same way. Any face can be approximated as a linear combination of the modes of variation:

$$\mathbf{m} = \bar{\mathbf{m}} + \sum_{i=1}^{j-1} \alpha_i \sigma_{m,i}^2 \mathbf{M}_i, \quad \mathbf{t} = \bar{\mathbf{t}} + \sum_{i=1}^{j-1} \beta_i \sigma_{t,i}^2 \mathbf{T}_i.$$

For convenience, we define the following vectors: $\mathbf{a} = [\alpha_1 \sigma_{m,1}^2 \ \dots \ \alpha_{j-1} \sigma_{m,j-1}^2]^T$ and $\mathbf{b} = [\beta_1 \sigma_{t,1}^2 \ \dots \ \beta_{j-1} \sigma_{t,j-1}^2]^T$ to model shape and texture respectively.

3 Shape Estimation

The majority of existing algorithms for fitting a morphable model utilise an analysis-by-synthesis approach. In this framework all parameters are estimated in a single non-convex cost function. To avoid the drawback of getting stuck in local minima, we decouple the estimation of shape and texture. Aldrian and Smith [AS] show that accurate shape parameters (comparable with state-of-the-art) can be obtained by simply fitting to sparse features in the image. In conjunction with shape parameters, the algorithm recovers geometric properties of the camera. The key advantage of this approach is that shape and texture estimation can both be posed as multilinear (and hence convex) fitting problems and solved independently.

We begin by estimating 3D shape parameters and the camera matrix by fitting to sparse feature points using the algorithm proposed in [AS]. This approach uses an empirical model of the generalisation capability of each feature point. This means errors can be appropriately penalised without having to tune a regularisation parameter to trade-off quality of fit against the likelihood of the solution. We assume the position of the sparse feature points is

known. The large body of work in facial feature tracking can be used to do this automatically. The estimated shape is used to construct the specular and diffuse spherical harmonic basis functions which are employed by the inverse rendering process.

4 Spherical Harmonic Lighting

Our starting point is the shape parameter, \mathbf{a} , which in turn determines a 3D mesh and per-vertex surface normals. In conjunction with the camera projection matrix we can test the visibility of each vertex and compute their projection to the 2D image plane. Hence, we have correspondence between some subset of the model vertices and the input image and the surface normal at the 2D positions of the projected vertices is known. We use a spherical harmonic representation to describe diffuse and specular reflectance.

Illumination and reflectance are estimated as follows. First, we estimate diffuse albedo and low frequency illumination. Second, we estimate low frequency specular reflectance such that the illumination is constrained to be consistent with the diffuse illumination. Finally, we estimate high frequency specular reflectance by solving an unconstrained optimisation problem.

Both Basri and Jacobs [10] and Ramamoorthi *et al.* [11] simultaneously and independently showed that illumination, bidirectional reflectance and irradiance can be expressed using spherical harmonics. Spherical harmonics form a set of orthonormal basis functions for the set of all square integrable functions defined on the unit sphere. They are functions of surface geometry and in polar coordinates they are defined as follows:

$$\mathcal{H}_{lm}(\theta, \phi) = \sqrt{\frac{2l+1}{4\pi} \frac{(l-m)!}{(l+m)!}} P_{lm}(\cos \theta) e^{lm\phi}.$$

Alternatively, they can be expressed in normalised cartesian coordinates (x, y, z) . The factor P_{lm} corresponds to the associated Legendre polynomial of order l and degree m . The following shows the expansion of an arbitrary, distant lighting function L in terms of spherical harmonics:

$$L(x, y, z) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \mathcal{L}_{lm} \mathcal{H}_{lm}(x, y, z). \quad (1)$$

The indices obey $l \geq 0$ and $-l \leq m \leq l$. The order, l , determines the polynomial degree used to compute the basis functions in terms of the surface normals. \mathcal{L}_{lm} are the corresponding weightings which are termed the lighting coefficients. The surface bidirectional reflectance distribution function (BRDF) can also be defined using an infinite series of spherical harmonic basis functions:

$$\hat{\rho} = 2\pi \int_0^{\pi/2} \sum_{l=0}^{\infty} \hat{\rho}_l \mathcal{H}_{l0}(\theta'_i) \mathcal{H}_{l0}(\theta'_i) \sin(\theta'_i) d\theta'_i. \quad (2)$$

Note, that BRDF is only a function of degree l , and that definition 2 only hold for symmetric BRDFs (see [11] for a detailed description). By combining 1 and 2, any symmetric reflectance function can be composed by multiplying corresponding frequency coefficients of lighting function and BRDF. In other words, the reflectance function is obtained by filtering the lighting function with the BRDF. The lighting function can empirically be determined using a chrome or mirror sphere placed in the scene and captured via a panoramic camera.

This is not feasible for most real-world scenarios. However, a low-pass filtered version of the reflectance function can be obtained from single photographs for Lambertian objects. Previous experimental results have shown that unconstrained complex illumination can well be approximated by a linear subspace, and a second degree approximation accounts for at least 98% in the variability of the reflectance function [9].

5 Image Formation Process

We assume a dichromatic reflectance model comprising additive Lambertian and specular terms. The specular function is assumed to be isotropic about the specular reflection direction but is otherwise unconstrained. Unlike [9, 10, 11], we allow any combination of directed, ambient or extended light sources.

Texture and Illumination To model unconstrained illumination of arbitrary colour, we construct a set of basis functions, $\mathcal{U}(x, y, z)$, derived from $\mathcal{H}(x, y, z)$. For a single vertex k , our modified basis functions are defined as follows:

$$\mathcal{U}(x, y, z)_k = \begin{bmatrix} \mathcal{H}(x, y, z)_k & \mathbf{0}^T & \mathbf{0}^T \\ \mathbf{0}^T & \mathcal{H}(x, y, z)_k & \mathbf{0}^T \\ \mathbf{0}^T & \mathbf{0}^T & \mathcal{H}(x, y, z)_k \end{bmatrix}. \quad (3)$$

We also take specular reflection of arbitrary unconstrained illumination into account. According to [9], a different set of basis function, $\mathcal{H}(x', y', z')_k$, can be constructed for that purpose by reflecting the normals about the viewing direction. We substitute the basis functions $\mathcal{H}(x, y, z)_k$ in equation 3 with the reflected ones and construct the specular set: $\mathcal{S}(x, y, z)_k$. We denote a particular column of this matrix as $\mathcal{S}(x, y, z)_k^c$, where $c \in \{r, g, b\}$. This notation will be used later when fitting the specular part. The full image formation process now takes the following form:

$$\mathbf{I}_{mod} = (\mathcal{U}\mathbf{I}) * (\mathbf{T}\mathbf{b} + \bar{\mathbf{t}}) + \mathcal{S}\mathbf{x},$$

where \mathbf{I} , \mathbf{b} and \mathbf{x} are the loadings for diffuse lighting, texture and specular contribution. The symbol $*$ in our image formation process denotes element wise multiplication.

Colour Transformation To make our fitting algorithm more flexible and allow fitting to different imaging properties, we estimate a linear colour transformation, $\mathbf{M}(\cdot) \in \mathbb{R}^{3 \times 3}$ and add offset, $\mathbf{o} \in \mathbb{R}^3$, firstly proposed in [9] to the illuminated estimations (\cdot) . Decomposition of \mathbf{M} into individual contributions is achieved as follows:

$$\mathbf{M} = \mathbf{G}\mathbf{C} = \begin{pmatrix} g_r & 0 & 0 \\ 0 & g_g & 0 \\ 0 & 0 & g_b \end{pmatrix} \cdot \left[c\mathbf{I} + (1 - c) \begin{pmatrix} 0.3 & 0.59 & 0.11 \\ 0.3 & 0.59 & 0.11 \\ 0.3 & 0.59 & 0.11 \end{pmatrix} \right],$$

where the entries g_r, g_g and g_b are gains for RGB respectively, and c corresponds to the contrast value. The gains are lower bound to 0 and contrast is constrained to lay between 0 and 1.

The Complete Model The entire image formation process is a multilinear system which consists of two nested bi-affine parts. For a single vertex, k , the image formation is modelled as:

$$\mathbf{I}_{mod,k} = \mathbf{M}[(\mathcal{U}_k\mathbf{I}) * (\mathbf{T}_k\mathbf{b} + \bar{\mathbf{t}}_k) + \mathcal{S}_k\mathbf{x}] + \mathbf{o}. \quad (4)$$

6 Inverse Rendering

We now show how the unknowns in equation 4 can be recovered. This amounts to a series of linear least squares problems. The entire system can be iterated to convergence but we have found that one pass, as described, is sufficient for good results. We begin by estimating the colour transformation parameters using the mean texture as initialisation and apply the inverse to the observations. We denote the colour corrected observations as $\bar{\mathbf{I}}$.

Diffuse Component The diffuse and specular shading coefficients, \mathbf{l} and \mathbf{x} , both depend on a single lighting function: $\mathcal{L} = [\mathcal{L}_r^T \mathcal{L}_g^T \mathcal{L}_b^T]$. As the lighting function can not be estimated directly from a single 2D image, we start by estimating \mathbf{l} and \mathbf{b} in a bilinear fashion. Ignoring the specular part at this stage, we minimise the following objective function, which depends on the colour transformation and observations:

$$E_d = \|\bar{\mathbf{I}} - (\mathcal{U}\mathbf{l}) * (\mathbf{T}\mathbf{b} + \bar{\mathbf{t}})\|^2. \quad (5)$$

Regularisation To prevent overfitting, we introduce two sets of prior on the parameters which encourage simplicity: $E_1 = \|\mathbf{b}\|^2$ and $E_2 = \|\mathbf{l}\|^2$. We also define a “grayworld” prior, E_3 , which prefers white illumination. We implement this constraint by encouraging the difference between \mathcal{L}_r^T , \mathcal{L}_g^T and \mathcal{L}_b^T to be small. The priors are added to equation 5 to form the overall cost function: $E_a = E_d + \lambda_1 E_1 + \lambda_2 E_2 + \lambda_3 E_3$.

Diffuse Bi-affine System The objective function E_a is convex in \mathbf{b} and \mathbf{l} . We treat both sets as independent contributions, find the partial derivatives, set to zero:

$$\frac{\partial E_a}{\partial \mathbf{b}} = \frac{\partial E_d}{\partial \mathbf{b}} + \frac{\partial \lambda_1 E_1}{\partial \mathbf{b}} = 0 \quad \text{and} \quad \frac{\partial E_a}{\partial \mathbf{l}} = \frac{\partial E_d}{\partial \mathbf{l}} + \frac{\partial \lambda_2 E_2}{\partial \mathbf{l}} + \frac{\partial \lambda_3 E_3}{\partial \mathbf{l}} = 0, \quad (6)$$

and solve for both sets using alternating-least-squares. In general the solution is independent of initialisation, although using the mean texture results in swift convergence within ≤ 5 iterations.

Specular Component The only remaining aspect of appearance to model is specular reflectance. This part is solved in two steps. For low frequency, $l \leq 2$, we use the illumination environment estimated in the diffuse fitting stage. For higher frequencies, $3 \leq l \leq 8$, we use an unconstrained optimisation procedure and hence the contribution of higher frequency illumination to specular reflectance is free to vary independently. The problem can be stated as: $\bar{\mathbf{I}}_s = \bar{\mathbf{I}}_{s,l} + \bar{\mathbf{I}}_{s,h}$. We also assume that specular reflectance is symmetric about the reflection vector.

Low Order Specular Reflectance Using the estimated parameters: \mathbf{b} and \mathbf{l} , we synthesise a diffuse-only image and subtract from the input image to obtain the specular-only image, $\bar{\mathbf{I}}_s$:

$$\bar{\mathbf{I}}_s = \bar{\mathbf{I}} - (\mathcal{U}\mathbf{l}) * (\mathbf{T}\mathbf{b} + \bar{\mathbf{t}}).$$

We clamp negative values to zero (these are caused by cast shadows or errors in the diffuse estimate). The lighting coefficients \mathcal{L}_{lm}^c are obtained by dividing diffuse coefficients, l_{lm} by the Lambertian BRDF parameters, which are constant for a given order.

Specular reflection requires an alternate basis-set constructed with respect to the reflection vector. Hence, we reflect the the viewing direction about the normals and define

new specular basis functions $\mathcal{S}(x', y', z')$. Since the diffuse illumination environment is already known, the isotropic specular reflectance function has only 3 free parameters $\hat{\tau}_l$, where $l \in \{0, 1, 2\}$ which can be obtained by solving the following linear system of equations:

$$\begin{aligned} \bar{\mathbf{I}}_{s,l} = & \hat{\tau}_0 \sum_c (\mathcal{S}_0^c \mathcal{L}_0^c) + \hat{\tau}_1 \sum_c (\mathcal{S}_{1-1}^c \mathcal{L}_{1-1}^c + \mathcal{S}_{10-10}^c \mathcal{L}_{10-10}^c + \mathcal{S}_{11-11}^c \mathcal{L}_{11-11}^c) + \\ & \hat{\tau}_2 \sum_c (\mathcal{S}_{2-2}^c \mathcal{L}_{2-2}^c + \mathcal{S}_{2-1}^c \mathcal{L}_{2-1}^c + \mathcal{S}_{20-20}^c \mathcal{L}_{20-20}^c + \mathcal{S}_{21-21}^c \mathcal{L}_{21-21}^c + \mathcal{S}_{22-22}^c \mathcal{L}_{22-22}^c). \end{aligned}$$

Multiplying specular BRDF parameters, $\hat{\tau}_l$, with the corresponding lighting coefficients, \mathcal{L}_{lm}^c , results in the specular coefficients, x_{lm}^c . These are concatenated to form vector $\mathbf{x}_l \in \mathbb{R}^{27}$.

Higher Order Specular Reflectance For orders $l \in \{3, \dots, 8\}$ both lighting and BRDF are unknown. In principle an exact and unique solution does not exist for this problem. However it is possible to solve an unconstrained problem and factor the solution. Again this can only be achieved up to a global scale factor. As this would not lead to significant advantage for our approach, we solve the unconstrained problem for higher order approximations:

$$\bar{\mathbf{I}}_{s,h}^c = [\mathcal{S}_{8-8}^c \quad \mathcal{S}_{8-7}^c \quad \dots \quad \mathcal{S}_{3-3}^c \quad \dots \quad \mathcal{S}_{33}^c \quad \dots \quad \mathcal{S}_{87}^c \quad \mathcal{S}_{88}^c] \mathbf{x}_h^c.$$

We solve for $\mathbf{x}_h^c \in \mathbb{R}^{54}$ using least-squares, and construct $\mathbf{x}_h = [\mathbf{x}_h^r \quad \mathbf{x}_h^g \quad \mathbf{x}_h^b] \in \mathbb{R}^{216}$.

Having an estimate for \mathbf{b}, \mathbf{l} and \mathbf{x} , we can synthesise vertex, k and solve for the second set of parameters. In the previous section, $\bar{\mathbf{I}}$, referred to the model without colour transformation applied to. We stick with this notation and denote $\bar{\mathbf{I}}_o$ for the model without offset applied to.

$$\mathbf{I}_{mod,k} = \mathbf{M} \bar{\mathbf{I}}_{o,k} + \mathbf{o} = \mathbf{G} \mathbf{C} \bar{\mathbf{I}}_{o,k} + \mathbf{o}.$$

This leads to the second bi-affine system:

$$\mathbf{E}_b = \|\bar{\mathbf{I}} - (\mathbf{G} \mathbf{C} \bar{\mathbf{I}}_{o,k} + \mathbf{o})\|^2.$$

Again, we set the partial derivatives with respect to the parameters \mathbf{g} and c to zero (similar to equation (6)), and solve using alternating-least-squares.

7 Experiments

We use the *Basel Face Model* [1] for our experiments. In addition to the training data, 10 out-of-sample faces in correspondence with the model are provided. Each face is rendered in 9 pose angles and 3 lighting conditions per pose. We conduct the following types of experiments to demonstrate the utility of our method. In the first experiment, we recover the 3D shape from sparse feature points of 2D renderings (provided with the model) using the method described in [1]. The shape and pose estimates are used to calculate spherical harmonic basis functions to recover diffuse and specular reflectance respectively. We also show how bilinear colour transformation can be used with the proposed method to fit to grayscale images. To demonstrate inverse-rendering under arbitrary illumination we use environment maps provided by the University of Southern California [2]. In the last experiment we show fitting results on 2D photographs taken with a Nikon D200 camera.

Synthetic Data We show qualitative and quantitative results for the 270 Basel renderings for texture only (shape estimates are provided by the method in [14]). We compare our results against the state-of-the-art *Multi feature fitting algorithm* [14]. Texture coefficients for that method are provided with the model release. We use the 60 most significant principal components to model texture. Our evaluation is based on the following two error measures:

$$E_g = \frac{1}{n} \|\mathbf{t}_g - \mathbf{t}_r\|^2 \quad E_m = \frac{\|\mathbf{t}_g - \mathbf{t}_r\|^2}{\|\mathbf{t}_m - \mathbf{t}_r\|^2}, \quad (7)$$

where, \mathbf{t}_g is the ground truth texture data provided with the out-of-sample faces and \mathbf{t}_m is the ground truth data projected into the 60 parameter model (i.e. the optimal model fit to the data). \mathbf{t}_r is the texture recovered using the proposed method. Individual values within each texture vector are within $\mathbb{R} \in [0, 1]$. E_m is a relative error measure which relates our result to the best possible the model can achieve for a particular out-of-sample face (a value of 1 denotes optimal performance).




Figure 1: Fitting results to out-of-sample renderings for two subjects. Left column shows ground truth (top) projected into the model using the 60 most significant modes (below). Second column: Top row shows out-of-sample renderings in different pose and illumination conditions. Second row shows texture reconstruction using to proposed method. Last row shows texture fitting results using the reference method [14].

Subject:		1	2	3	4	5	6	7	8	9	10	Mean
$E_g \times 10^{-3}$	Prop.	6.0	4.1	8.2	3.7	17.0	3.5	5.0	4.6	5.3	4.7	6.2
	[14]	8.8	7.5	7.5	4.4	14.2	11.0	12.4	12.0	7.4	5.2	9.0
E_m	Prop.	5.0	3.1	6.3	4.3	12.9	3.2	5.3	4.7	8.1	4.3	5.7
	[14]	7.4	5.7	5.7	5.1	10.8	9.8	13.3	12.2	11.4	4.7	8.6

We achieve an average error ratio, \bar{E}_m , of 5.69 compared to 8.61 for the reference method. The absolute reconstruction error compared to ground truth is , \bar{E}_g , is 6.2×10^{-3} for the proposed method versus 9.0×10^{-3} for the reference method. The above table shows quantitative results for the 10 subjects, averaged over all pose angles and illumination conditions.

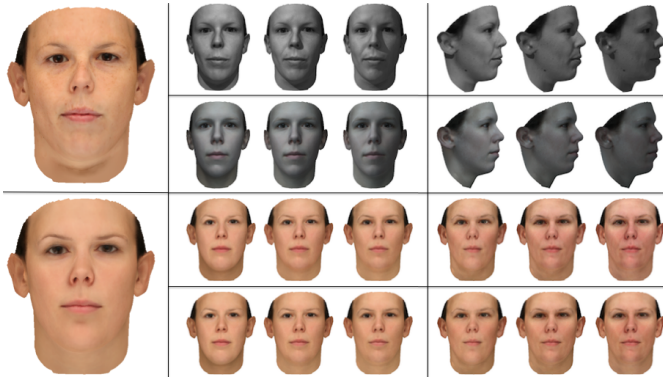


Figure 2: Fitting results to grayscale images. First column shows ground truth shape and texture (top) and texture projected into the model (60 modes). Column 2 and 3: Top row shows input images 0 and -70 in 3 illumination conditions respectively. Second row shows fitting results with all parameters estimated, including shape. Third row shows the fitting results for texture only. Last row shows fitting results to the equivalent RGB images using the proposed method. They demonstrate how fitting accuracy to grayscale images degrades for varying illumination in more extreme pose angles.




Figure 3: We use environment maps (top left) and calculate diffuse and specular approximations (below). These are used to render an out-of-sample face (top row of right hand side. The image furthest rights shows ground truth texture). Below are the fitting results using the proposed method. The last row on the left, shows the re-estimated environment maps using parameters obtained in the fitting process.

In figure 1 we show diffuse albedo estimates for 2 subjects in 5 pose angles. Illumination condition is shown randomly for each pose.

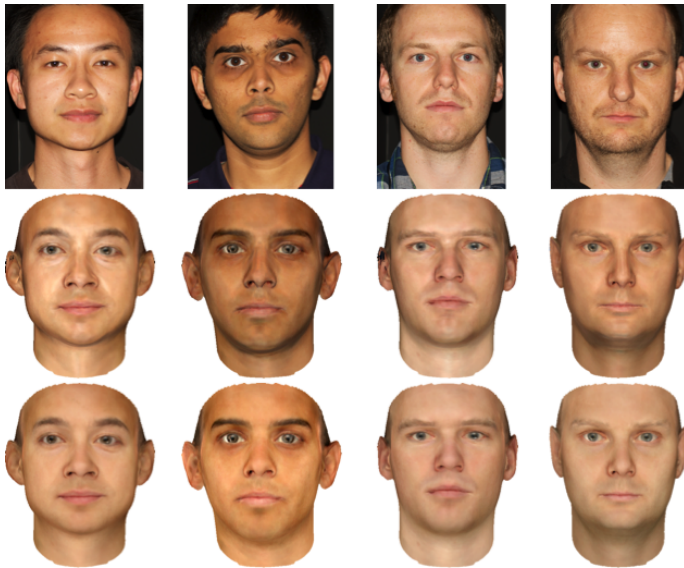


Figure 4: Face shape and appearance modelling from photographs. Top row shown photographs of 4 subjects, not part of the training data. Second row shows a full model of the photograph (Shape, camera properties, albedo, colour transformation, diffuse and specular lighting). Last row shows diffuse albedo only.

Grayscale Synthetic Data Our method can be used to fit the colour model to grayscale imagery. We convert the rendering for one subject to grayscale and use the colour transformation described in section 5 to fit. Figure 2 shows the results for subject one of the test-set in two poses, 0 and -70° , and 3 illumination conditions. For comparison, we also show fitting results for the original colour images. Using the error measure introduced in equation 7, and averaging over all 27 renderings we obtain the following values for face one: $E_g = 6.3$ and $E_m = 5.27$.

Real World Illumination We used high dynamic range environment maps provided by the University of Southern California [6] to render images of the out-of-sample faces under complex illumination. In figure 3, we demonstrate how the proposed method can be used to accurately and efficiently deconvolve, diffuse lighting, specular reflectance and texture.

Real World Images We captured photographs of 4 subjects disjoint from the training data using a Nikon D200 camera. This last part of the experimental section demonstrates the full inverse rendering process, including shape and camera properties, see figure 4.

8 Conclusions

We have presented a practical and very general approach to fitting a morphable model to face images. Our evaluation shows that the method performs well, even under highly complex illumination. In future, we will evaluate how the nature of the illumination environment affects the accuracy of the estimated texture. We will also explore how the residual fitting

error can be used to further refine shape and texture estimates by moving outside the span of the global statistical model.

References

- [1] Oswald Aldrian and William A. P. Smith. A linear approach of 3d face shape and texture recovery using a 3d morphable model. In *Proceedings of the British Machine Vision Conference*, pages 75.1–75.10. BMVA Press, 2010. ISBN 1-901725-40-5.
- [2] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(2):218–233, 2003.
- [3] V. Blanz and T. Vetter. A morphable model for the synthesis of 3D faces. In *Proc. SIGGRAPH*, pages 187–194, 1999.
- [4] V. Blanz, A. Mehl, T. Vetter, and H-P. Seidel. A statistical method for robust 3d surface reconstruction from sparse data. In *Proc. 3DPVT*, pages 293–300, 2004.
- [5] C. Donner, T. Weyrich, E. d’Eon, R. Ramamoorthi, and S. Rusinkiewicz. A layered, heterogeneous reflectance model for acquiring and rendering human skin. *ACM Trans. Graphic. (Proceedings of SIGGRAPH Asia)*, 27(5), 2008.
- [6] University of Southern California. High-resolution light probe image gallery, 2011. <http://gl.ict.usc.edu/Data/HighResProbes>.
- [7] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter. A 3D face model for pose and illumination invariant face recognition. In *Proc. IEEE Intl. Conf. on Advanced Video and Signal based Surveillance*, 2009.
- [8] Ravi Ramamoorthi. Modeling illumination variation with spherical harmonics. In *Face Processing: Advanced Modeling and Methods*. Academic Press, 2005.
- [9] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for reflection. *ACM Trans. Graph.*, 23(4):1004–1042, 2004. ISSN 0730-0301. doi: <http://doi.acm.org/10.1145/1027411.1027416>.
- [10] S. Romdhani and T. Vetter. Estimating 3D shape and texture using pixel intensity, edges, specular highlights, texture constraints and a prior. In *Proc. CVPR*, volume 2, pages 986–993, 2005.
- [11] Sami Romdhani, Volker Blanz, and Thomas Vetter. Face identification by fitting a 3d morphable model using linear shape and texture error functions. In *European Conference on Computer Vision*, pages 3–19, 2002.
- [12] Hyunjung Shim, Inwoo Ha, Taehyun Rhee, James Dokyoon Kim, and Changyeon Kim. A probabilistic approach to realistic face synthesis. In *Proc. ICIP*, 2010.
- [13] L. Zhang and D. Samaras. Face recognition from a single training image under arbitrary unknown lighting using spherical harmonics. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(3):351–363, 2006.