

Efficient Global Illumination for Morphable Models

Andreas Schneider, Sandro Schönborn, Bernhard Egger, Lavrenti Frobeen, Thomas Vetter
University of Basel, Switzerland

{andreas.schneider, sandro.schoenborn, bernhard.egger, l.frobeen, thomas.vetter}@unibas.ch

Abstract

We propose an efficient self-shadowing illumination model for Morphable Models. Simulating self-shadowing with ray casting is computationally expensive which makes them impractical in Analysis-by-Synthesis methods for object reconstruction from single images. Therefore, we propose to learn self-shadowing for Morphable Model parameters directly with a linear model. Radiance transfer functions are a powerful way to represent self-shadowing used within the precomputed radiance transfer framework (PRT). We build on PRT to render deforming objects with self-shadowing at interactive frame rates. It can be illuminated efficiently by environment maps represented with spherical harmonics. The result is an efficient global illumination method for Morphable Models, exploiting an approximated radiance transfer. We apply the method to fitting Morphable Model parameters to a single image of a face and demonstrate that considering self-shadowing improves shape reconstruction.

1. Introduction

Morphable Models [6] describe the variation of shape and texture of objects such as faces. They allow to draw new samples, create animations of morphing between two instances, and to reconstruct shape, texture, and illumination from a single image by adapting model parameters to it. Currently Morphable Models are oblivious to self-shadowing. Self-shadowing has a large effect on images. The extent depends on how much the shape shadows itself and under which illumination it is lit. In faces, expressions and aging are deformations resulting in substantial self-shadowing. Self-shadowing comes with a significant computational cost which makes morphing and model adaptation impractical. Being able to interact with the Morphable Model or making inverse rendering practical requires a fast self-shadowing model. In this work, we propose a method for fast and realistic self-shadowing for the Morphable Model.

Simulating self-shadowing requires tedious sampling of

the hemisphere for every point on the surface by casting rays [17]. We represent self-shadowing with radiance transfer functions. For a given shape PRT [20] allows to precompute them. The shape can then be lit under arbitrary low-frequency environment maps [20], and BRDFs [8] at interactive frame rates. This method is suitable for rigid objects. PRT is expensive and unfeasible for rendering deforming objects like those arising from Morphable Models. For every surface change we would have to compute the radiance transfer with ray-casting.

In the Morphable Model, shape is modelled linearly. Similar shapes are close to each other in the shape space. It can be expected that self-shadowing of neighbouring shapes is similar too. As a consequence, we compute self-shadowing only for a few hundred shapes instead of expensive ray casting for every possible shape. In this work we model the radiance transfer field of the Morphable Model. The idea is to approximate radiance transfer directly from shape parameters. We learn the mapping between Morphable Model shape parameters with linear regression. We assume the object to be of diffuse reflectance to be lit by a low-frequency environment map represented with spherical harmonics. Extension to other BRDFs and interreflections is possible. In the methods section we describe the linear approximation of the radiance transfer field of the Morphable Model. In the experiments section we show that a linear approximation is sufficient and compare the proposed fast Morphable Model self-shadowing model (MoMo-PRT) to the exact and expensive PRT method. We show that reconstructing face shape from a single image is improved with the proposed method. Figure 1 gives an overview over the proposed approach.

1.1. Related Work

Morphable Models were first introduced by Blanz and Vetter [6]. They model face shape and texture independently of each other. Face shape and texture are modelled as a linear combination of Principal Component Analysis (PCA) basis vectors computed from exemplar 3D face scans. The Analysis-by-Synthesis method of Schönborn *et al.* [18] uses a MCMC Metropolis Hastings method to re-

construct shape, texture, and illumination from a single image. Analysis-by-Synthesis methods require realistic and fast rendering methods.

Screen-space methods are fast in calculating effects like ambient occlusion [22] frame by frame. In screen space directed occlusion [15] the local contribution of self-shadowing is computed in screen space and the global contribution on the coarsened mesh by ray tracing. Screen-space methods do not consider the distribution of shadows created by the Morphable Model shape distribution. Our proposed method exploits the Morphable Model shape statistic to create shadowing. Many Morphable Model adaptation frameworks exploit the deterministic connection between shadowing and shape parameters to estimate shape, texture, and illumination. In screen space methods the connection between self-shadowing and shape parameters is difficult to establish. MoMo-PRT provides trivial derivatives of self-shadowing with respect to shape parameters which would be difficult with screen-space methods.

This work builds on PRT, which represents self-shadowing with linear radiance transfer functions. It was introduced by Sloan and Kautz [20] and speeds up the rendering of non-deforming objects with global illumination. They use the concept of radiance transfer functions to represent self-shadowing. They compute radiance transfer for every point on the surface by simulating light-surface interaction, once, for the given shape with ray tracing. Because the radiance transfer of the shape can be reused for different illuminations, interactive rendering of non-deforming surfaces but changing environment map illumination becomes possible in real-time.

Early Morphable Model based image analysis methods included self-shadowing using a single light source and a Phong illumination model ([6], [16]). Aldrian and Smith *et al.* [2] and Schönborn *et al.* [18] use spherical harmonics to approximate the environment map and can produce more realistic illumination settings but lack shadowing. The work of Shahlaei and Blanz [19] uses an illumination cone to combine multiple light sources to approximate the environment map and includes cast shadows in rendering. An approximation to self-shadowing with ambient occlusion and bent normals [11] for Morphable Models was proposed by Aldrian *et al.* [1]. They show that texture and illumination reconstruction can be improved with their approach. These methods are coarse approximations of true self-shadowing. Radiance transfer functions however, model the full surface-light interaction. Self-shadowing expressed within the PRT framework results in more realistic rendering.

Shape from shading methods as post processing of Morphable Model reconstructions are used by different methods ([21, 14]) to obtain more detailed reconstructions. These methods do not respect self-shadowing and are therefore not

suited to reconstruct details lit by self-shadowing.

Previous work coupling shape and ambient occlusion was done by Kontkanen and Aila [9] who model ambient occlusion of an animated character as a linear combination of animation parameters. In this work we show how to model global illumination with radiance transfer matrices for Morphable Models illuminated by low frequency environment maps.

2. Methods

For efficient rendering we directly estimate self-shadowing from Morphable Model shape parameters using linear regression. We can thus eliminate the computationally expensive ray casting calculation during rendering. The overall computational burden of computing self-shadowing for a deforming object is significantly reduced. Calculation of a radiance transfer matrix, as required by PRT rendering, is reduced to a matrix-vector multiplication. For full benefit of explicit self-shadowing, we correct texture images and remove effects of ambient occlusion. Finally, we also present how the efficient, linear self-shadowing model can be applied to efficiently estimate an unknown illumination environment map from a single face image. Figure 1 gives an overview over the methods.

2.1. Morphable Model

Our 3D Morphable Face Model is derived from the the Basel Face Model [13], built from face scans of 200 people taken with a structured light 3D scanner. We extended it with a multi-linear model to incorporate facial expressions as proposed by Amberg *et al.* [3]. The model is composed of two independent PCA models for shape and texture as well as an additional PCA model for facial expressions. Expressions are modelled as difference to the neutral faces.

2.2. Morphable Model Radiance Transfer

For Morphable Model Radiance Transfer (MoMo-PRT), we estimate self-shadowing radiance transfer directly from the model’s shape representation. The method is based on PRT [20] but introduces a computationally efficient way of finding transfer matrices.

Within the PRT framework, emitted radiance r on the surface of an illuminated object is calculated by linear transfer of incoming irradiance l . All functions on the sphere are expanded in a Spherical Harmonics basis. For every point i on the surface this is

$$\mathbf{r}_i = \mathbf{T}_i \mathbf{l}. \quad (1)$$

The transfer matrix \mathbf{T}_i describes the surface-light interaction and includes all self-shadowing effects. In PRT, it needs to be calculated by solving a full ray casting integration for each object while viewpoint and illumination changes do not require recomputing.

Morphable Model - Precomputed Radiance Transfer (MoMo-PRT)

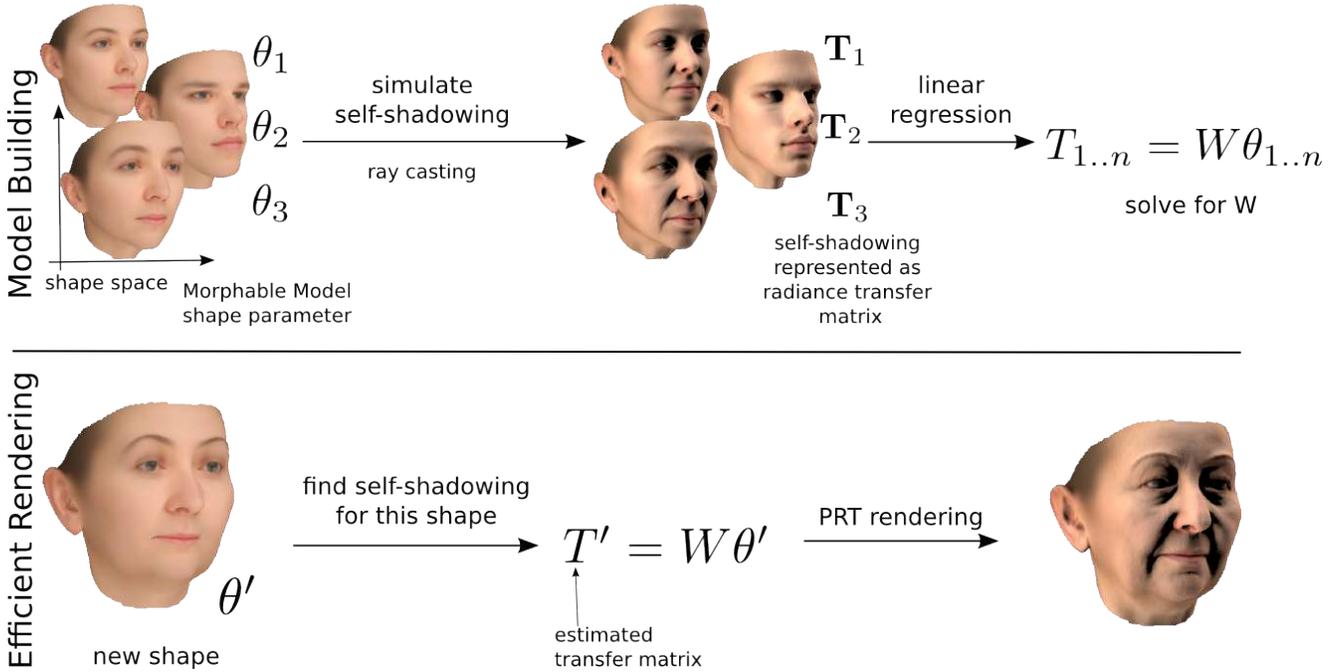


Figure 1: Efficient self-shadowing for the Morphable Model (MoMo-PRT). We approximate self-shadowing for the Morphable Model parameters. **Model building:** First, we simulate self-shadowing for different faces in face space. Second, we find a linear mapping between shape parameters and self-shadowing parameters. **Efficient Rendering:** For a new Morphable Model shape, we estimate its corresponding self-shadowing parameters using the linear model. This allows to efficiently compute a rendering with self-shadowing for arbitrary Morphable Model parameters.

Observed radiance L_i at a position i is a linear combination of radiance \mathbf{r} with spherical harmonics basis functions \mathbf{Y} evaluated in normal direction \mathbf{n} of the surface

$$L_i = (\mathbf{Y}(\mathbf{n}_i) * \mathbf{b}_i)^t \mathbf{r}_i = (\mathbf{Y}(\mathbf{n}_i) * \mathbf{b}_i)^t \mathbf{T}_i \mathbf{l}. \quad (2)$$

$\mathbf{Y}(\mathbf{n}_i) * \mathbf{b}_i$ is the convolution of spherical harmonics basis functions with the BRDF \mathbf{b} . In this work, we assume a constant BRDF (Lambertian reflectance) on the whole surface and expand functions on the sphere using three bands of real spherical harmonics $Y_{lm}, l = 0, 1, 2$ (9 coefficients).

MoMo-PRT. For a Morphable Model, we model a linear relationship $\mathbf{W}_i(\theta)$ between radiance transfer matrices and shape parameters θ at every vertex i of the mesh

$$\mathbf{T}_i = \mathbf{W}_i(\theta) \quad (3)$$

and therefore the outgoing radiance

$$\mathbf{r}_i = \mathbf{W}_i(\theta) \mathbf{l}. \quad (4)$$

Matrices within triangles are found by barycentric interpolation.

Training. To actually find mappings \mathbf{W} , we perform multivariate linear regression on a training data set consisting of pairs of shape parameters and corresponding transfer matrices \mathbf{T} . Transfer matrices for training are calculated by full self-shadowing integration using ray tracing for each example shape on each vertex of the mesh.

Example shapes are obtained by varying principal components of the Morphable Model. We sample components at ± 2 standard deviations. More specifically, samples lie on a hypercube in parameter space with edge length 4. We systematically sample all corners on the hypercube for the first three principal components of shape and expression.

For every vertex i , we search the linear transform \mathbf{W}_i

$$\mathbf{W}_i \theta + \tilde{\mathbf{T}}_i^0 = \tilde{\mathbf{T}}_i. \quad (5)$$

Here, \mathbf{W}_i is a matrix mapping shape parameters to the vectorized transfer matrix $\tilde{\mathbf{T}}_i$. Below, we treat the constant offset $\tilde{\mathbf{T}}_i^0$ implicitly and remove it from notation to prevent clutter.

For each vertex i , we get a system of linear equations for n training examples

$$\mathbf{W}_i \theta = [\tilde{\mathbf{T}}_{i_1}, \tilde{\mathbf{T}}_{i_2}, \dots, \tilde{\mathbf{T}}_{i_n}]. \quad (6)$$

We render $n = 600$ training examples and have a dimensionality $9 \times 9 = 81$ of $\tilde{\mathbf{T}}$. We then obtain \mathbf{W}_i for each vertex as the least-squares solutions of above systems.

Shelf-shadowing approximated by the proposed linear method is visually indistinguishable from the full PRT computation while dramatically faster, see Fig. 2 and Table 1.

2.3. Self-Occlusion Correction of Textures

Textures of face scans contain shading due to self-shadowing even when lit with ambient illumination. In the classical Morphable Model, the texture model still contains such shading. Because shape and texture are modelled independently, shading due to self-occlusion can be expressed with the texture model, without changing the shape. With explicit shadow rendering, we correct scan textures to prevent double application of self-shadowing effects.

We remove the shading due to ambient occlusion similarly to the approach of Beeler *et al.* [5]. To remove shading due to self-occlusion, we subtract it from captured textures under ambient illumination. We explicitly calculate shelf-shadowing for the 3D scan geometry and find the amount of shading d due to self-shadowing by subtracting a rendering without self-shadowing $R_{\bar{s}}$ from a rendering with self-shadowing R_s . The corrected texture image $I_{\bar{s}}$ is obtained by subtracting the shading due to self-shadowing d from the input image I_s .

$$I_{\bar{s}} = I_s - (R_s - R_{\bar{s}}). \quad (7)$$

Physically-based rendering relies on a linear color space. We therefore linearize input colors in the sRGB color space with gamma correction before subtraction and rendering. The texture model is also built in linear color space.

2.4. Model Fit to an Image

The Morphable Model provides rich prior knowledge to obtain a 3D reconstruction of a face from a single image in an Analysis-by-Synthesis manner. Fitting the model to the image reconstructs shape, texture, pose and illumination. Our linear shadowing model can be directly applied to inverse rendering in the fitting process. It is possible to efficiently obtain the illumination condition of the face by solving a linear system which relates observed pixel colors, surface normals, self-occlusion, albedo and the environment map. Such estimation of illumination respects self-shadowing as an effect of surface-light interaction. For adapting the model to an image, we employ the Markov chain Monte Carlo fitting framework of Schönborn *et al.* [18].

To obtain the illumination estimation, we rewrite (2) as a system of linear equations for N observed radiance values $\mathbf{L} = [L_1 \ L_2 \ \dots \ L_N]^T$

$$\mathbf{L} = \mathbf{A}\mathbf{l}. \quad (8)$$

The rows of A are

$$\mathbf{t}_i = (\mathbf{Y}(\mathbf{n}_i) * \mathbf{b}_i)^t \mathbf{T}_i. \quad (9)$$

We use our linear estimate of the transfer matrices $\mathbf{T}_i = \mathbf{W}_i \theta$ in above system and obtain the environment map \mathbf{l} as the least-squares solution.

3. Experiments

We evaluate the approximation quality of the Morphable Model with self-shadowing and compare its 3D reconstruction performance to a Morphable Model without self-shadowing. The model with linear approximate shelf-shadowing can improve on face reconstruction on synthetic and real-world data.

3.1. Approximation Quality

We quantitatively compare renderings obtained with MoMo-PRT to its ground truth, a full PRT shelf-shadowing calculation. We evaluate the image difference between the obtained face renderings as a measure of error. We choose all other settings identical and directly compare the effects of including shelf-shadowing. To get a diverse, averaged result, we randomly sample face instances at increasing distance from the mean and thus decreasing statistical plausibility. For each distance value, we draw 10 samples, render the local illumination model, our linear approximation with MoMo-PRT and do the full shadowing calculation of PRT. The renderings use two different illumination setups, an evenly-lit ambient setup and a shadowing scene with illumination from the side. Both renderings are compared to full PRT calculation. The image reconstruction error is a pixel-wise Root Mean Square (RMS) difference of RGB values. For comparability, it is scaled relative to the maximally possible value (white). The results are presented in Fig. 3.

While the local model shows a difference of 3% and higher, the linear approximation is consistently below 1.5%, even below 0.5% for ambient light (only ambient occlusion effect). We observe a slightly increasing reconstruction error for more extreme shapes (farther from mean), which is completely expected. However, the MoMo-PRT deals with it quite well and only shows a very moderate increase for plausible shapes.

The low reconstruction error of the proposed method corresponds to a perceptually hardly visible difference. The

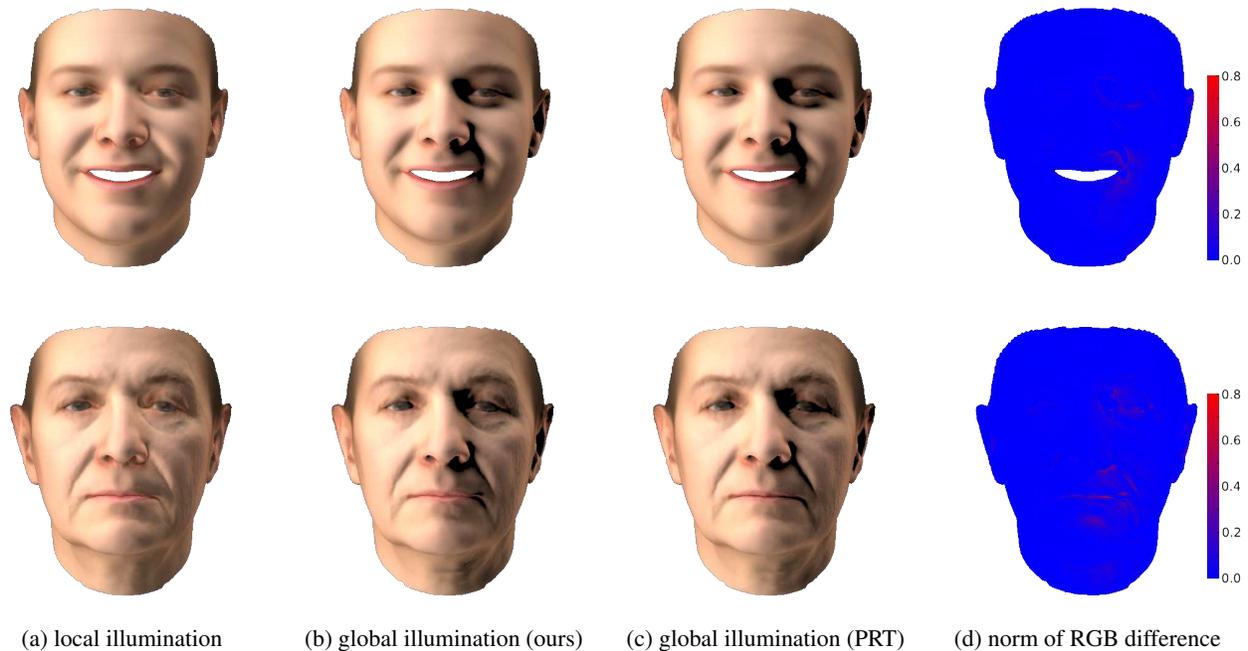


Figure 2: A comparison of the different rendering methods with the same environment map applied to a face with expression (**top**) and an aged face (**bottom**). The local model assumes Lambertian reflectance and contains no global illumination effects (**a**). The proposed model consists of local illumination plus efficient self-shadowing (**b**). It approximates physically based, accurate but slow self-shadowing of the PRT method (**c**). The approximate MoMo-PRT method (**b**) is visually indistinguishable from the full calculation (**c**). In contrast, the local model misses shadows to the left of the nose and cannot render the nasolabial fold with a satisfying impression of depth. This can be seen in (**d**). It shows the norm of the RGB difference vectors between (**b**) and (**c**). The mean error for the top row is 0.01 and for the bottom it is 0.02. The maximal possible deviation is 1.73. The maximum deviation is 0.35 and 0.67.

proposed model leads to renderings of faces with expression and aged faces with deeper folds which are visually indistinguishable from the full PRT calculation, see Fig. 2.

The small approximation error shows that linear methods are justified for this application. We do not see a benefit in the additional precision we would gain with non-linear methods due to the added cost. At least within the range of the shape variations we apply the model to.

3.2. Computational Performance Comparison

We compare the speed of rendering different faces with the proposed method to the full self-shadowing calculation of PRT. We sample random faces from the Morphable Model and render them with the simple local illumination model, the proposed model and with full PRT. We timed rendering and drawing a model sample in all three models. All meshes consist of roughly 50 000 vertices and 100 000 triangles. For a fast triangle lookup in the full PRT integration [20], we use an octree [12] implementation. It is to note that we are using a software renderer for all methods. To

measure time of computation we repeated the experiment ten times on an Intel Xeon CPU with 4 cores (8 threads at 3.6Hz). One thread each was used for the local and MoMo-PRT models and six for the full PRT simulation. In Table 1, we present the dramatic speedup of roughly $685\times$ using the linear approximation MoMo-PRT compared to the full self-shadowing simulation. For comparison, the local illumination model is still faster but does not consider any shadowing effects and generates less realistic face renderings. Whilst the full PRT rendering is not suitable for real-time-applications the proposed MoMo-PRT rendering can be performed in real-time.

3.3. Reconstruction Experiment on Artificial Data

We evaluate the self-shadowing model in terms of its quality of face reconstruction from images in a model fitting application. We therefore compare Morphable Model face reconstruction of shape, texture and illumination on artificial renderings to their ground truth values. The artificial target images depict a single face, illuminated by shadowing

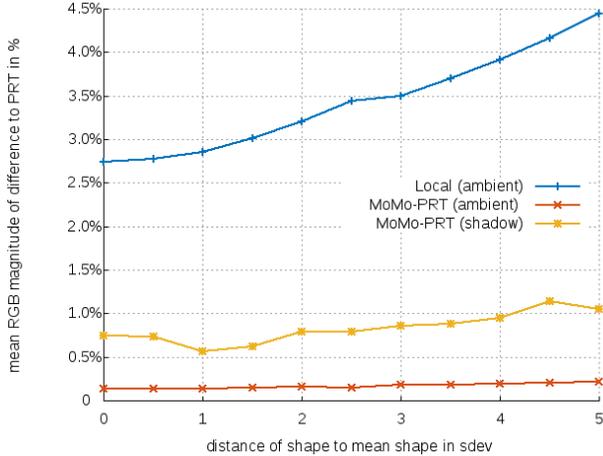


Figure 3: **Gap between different models and true self-shadowing.** Comparison of the local illumination model and MoMo-PRT with true self-shadowing for random face shapes with increasing distance to the mean. We compare renderings using the respective approximate model (**MoMo-PRT** and **local illumination**) with true self-shadowing under two different illumination conditions, ambient and with strong self-shadowing under illumination from the side (**shadow**). The approximation quality of MoMo-PRT is very high. It reproduces the image obtained by full self-shadowing calculation with less than 1.5% difference. However, local illumination cannot capture shadowing and reproduces the image considerably worse. Difference values are relative average RGB magnitudes of pixel-wise differences.

Table 1: Performance comparison. All values are relative to the full simulation. The mesh consists of 50 000 vertices and 100 000 triangles. Besides rendering, the measured times contain drawing a model sample in all three methods. More details can be found in Section 3.2.

Model	Speedup
PRT	1
MoMo-PRT	685
local	1 410

light, as used in Fig. 2. Model fitting is based on the MCMC method presented in [18]. We use their presented setup but replace illumination estimation with our own shadow-aware version presented in Section 2.4.

We render 20 randomly generated face images using a full physically based self-shadowing simulation as target images. Reconstructions are obtained by fitting the Mor-

phable Model with our linear radiance transfer extension and with a local illumination model only. The fit is properly initialized at the right face pose to focus on actual face reconstruction rather than initialization. We evaluate fits on the best of 10 000 drawn samples.

Evaluation is based on RMS differences to the ground truth data used for rendering. We evaluate shape difference, albedo reconstruction, illumination estimation and image reconstruction. Shape is a point-wise comparison of the 3D shape positions at each vertex. Texture is compared directly on a normalized image without any illumination. Illumination is compared by rendering a sphere under the estimated environment map and comparing the resulting images [4]. The image reconstruction error measures how well we can approximate the target image with the model fit. All color values are RMS differences in RGB space, see above. Shape difference is in mm. We test if improvements by the proposed method over the baseline are statistically significant with the Wilcoxon signed-rank test on the sample of 20 examples.

The experiment demonstrates a significant superior shape reconstruction performance when self-shadowing is considered. Also, the final model-reconstructed image is more similar to the original image. The effect can be emphasized for shapes which show more geometric details, such as aged faces. Table 2 presents a summary of the reconstruction results.

Table 2: Reconstruction performance of the proposed self-shadowing model (**global**) and the model with only local illumination (**local**) when adapted to an image. Reconstruction of shape and image are significantly (*) better with self-shadowing. Reconstruction errors measure Root Mean Square (RMS) differences between reconstruction and ground truth, for details refer to Section 3.3.

	local	global (ours)
shape [mm]	2.51	2.12*
texture	0.11	0.11
illumination	0.14	0.14
image	0.14	0.12*

3.4. Reconstruction Experiment on Photographs

To offer more intuitive analysis of fitting performance, we qualitatively demonstrate the result also on a few real-world images. We manually initialized fitting with user-provided landmarks locations. Fitting is performed with the same settings as in the above experiment. In Fig. 4 we present the resulting image reconstructions of our Morphable Model with self-shadowing and compare it to a fit obtained with a locally illuminated model.

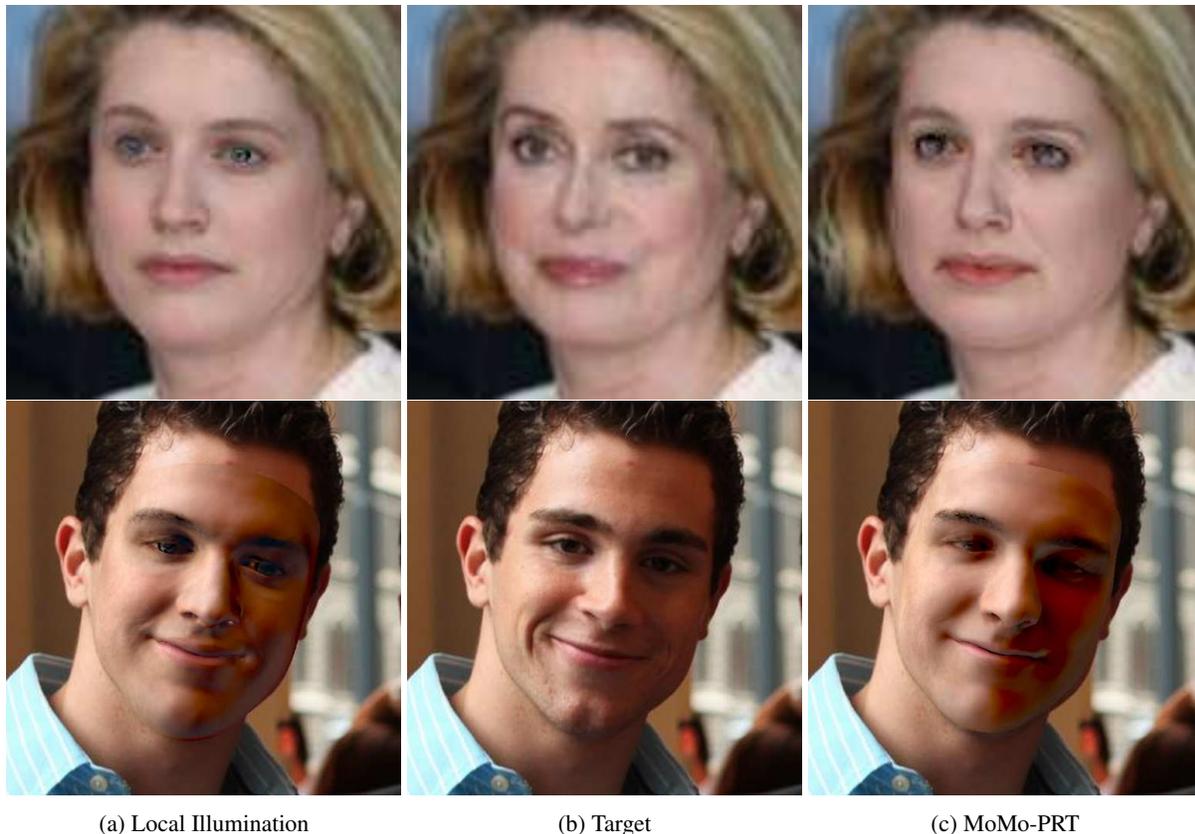


Figure 4: **Morphable Model** fits to target photograph (b), with the local illumination model (a) and our MoMo-PRT method (c). Improved aging and nasolabial fold in the first row. Second row exhibits self-shadowing: the left eye and left part of the mouth is less visible because of self-shadowing. Images are from Labeled Faces in the Wild (LFW) ([7]) and Annotated Facial Landmarks in the Wild (AFLW) ([10])

We performed a more detailed analysis of how the fitting result is composed of the different components of the face model. Whilst the reconstruction of the local model often looks feasible too, it tends to explain illumination effects by the texture model (see Fig. 5).

The resulting fits demonstrate the increased quality of reconstruction. With characteristic shapes and shadowing illuminations we observe the largest difference in image reconstruction between ours and the local model. The results of the self-shadowing based image analysis appear better in regions where shadows have strong effects on facial appearance. Mainly in the eye, nose and nasolabial fold regions the image reconstruction is better and those regions are better explained by shape and illumination and not the texture model. Nasolabial folds tend to be more pronounced and shadows can be reproduced better.

4. Conclusion

We presented an efficient global illumination model that allows interactive rendering of deformable Morphable Model instances with self-shadowing. We combined Morphable Models with Precomputed Radiance Transfer by modeling radiance transfer as a linear function of shape. We thus extended the applicability of efficient PRT rendering to deformable shapes. For the Morphable Model of faces, we achieved a visually indistinguishable approximation to the expensive ray casting calculating of self-shadowing but at a far lower computational load. Interactive rendering speed and a linear relation make inverse rendering including self-shadowing from a single image practical. We demonstrated that fitting a Morphable Model with self-shadowing to images improves face reconstruction and respects shadowing, for synthetic and real-world applications. The linear approach is open to integrate inter-reflections, arbitrary BRDFs and subsurface scattering in the future. For future work non-linear methods could be explored to improve the

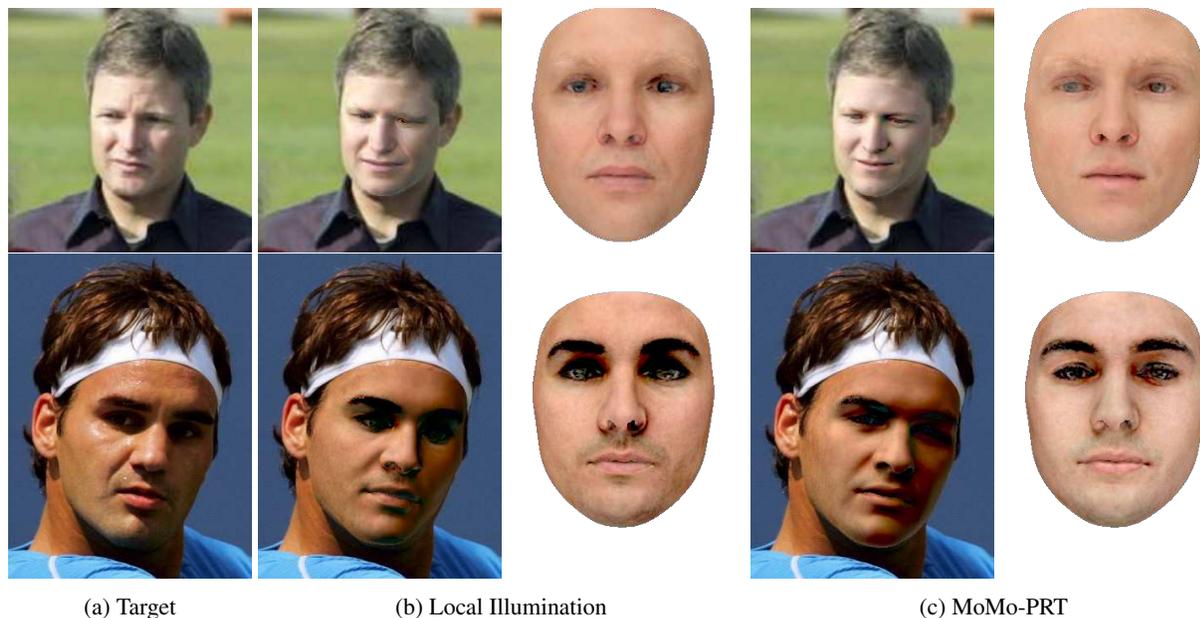


Figure 5: **Explanation of target image (a) with either texture or self-shadowing.** Reconstruction with local model (b). Reconstruction with our model (c). Both reconstructions hold a feasible image explanation (left). The constructed texture rendered on mean shape lit with ambient illumination (right) reveals the how the different models explain the target image. For the first example (top) the nasolabial fold is explained in the texture using the local illumination model and with shape and illumination using the proposed MoMo-PRT. The second example depicts the dark eye region in the target image reconstructed by the local illumination model explaining shadowing in the texture by making it unnaturally darker around the eyes. Our model explains the shadowing of the eye with self-shadowing. Upper Image: LFW [7], lower Image: KEYSTONE/EPA/Justin Lane

radiance transfer regression from shape parameters. To improve quality and speed MoMo-PRT could be combined with screen-space methods. Derivatives of self-shadowing with respect to shape parameters can be computed simply which opens the approach to other methods of shape reconstruction and possibly other applications in computer vision.

References

- [1] O. Aldrian and W. A. Smith. Inverse rendering of faces on a cloudy day. In *European Conference on Computer Vision*, pages 201–214. Springer, 2012. 2
- [2] O. Aldrian and W. A. Smith. Inverse rendering of faces with a 3d morphable model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(5):1080–1093, May 2013. 2
- [3] B. Amberg, R. Knothe, and T. Vetter. Expression invariant 3d face recognition with a morphable model. In *Automatic Face & Gesture Recognition, 2008. FG'08. 8th IEEE International Conference on*, pages 1–6. IEEE, 2008. 2
- [4] J. T. Barron and J. Malik. Shape, illumination, and reflectance from shading. *IEEE transactions on pattern analysis and machine intelligence*, 37(8):1670–1687, 2015. 6
- [5] T. Beeler, D. Bradley, H. Zimmer, and M. Gross. Improved reconstruction of deforming surfaces by cancelling ambient occlusion. In *European Conference on Computer Vision*, pages 30–43. Springer, 2012. 4
- [6] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 187–194. ACM Press/Addison-Wesley Publishing Co., 1999. 1, 2
- [7] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller. Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Technical Report 07-49, University of Massachusetts, Amherst, October 2007. 7, 8
- [8] J. Kautz, J. Snyder, and P.-P. J. Sloan. Fast arbitrary brdf shading for low-frequency lighting using spherical harmonics. *Rendering Techniques*, 2:291–296, 2002. 1
- [9] J. Kontkanen and T. Aila. Ambient occlusion for animated characters. In *Rendering techniques*, pages 343–348. Cite-seer, 2006. 2
- [10] M. Köstinger, P. Wohlhart, P. M. Roth, and H. Bischof. Annotated facial landmarks in the wild: A large-scale, real-world database for facial landmark localization. In *Computer*

Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on, pages 2144–2151, 2011. [7](#)

- [11] H. Landis. Production-ready global illumination. *Siggraph course notes*, 16(2002):11, 2002. [2](#)
- [12] D. Meagher. Geometric modeling using octree encoding. *Computer graphics and image processing*, 19(2):129–147, 1982. [5](#)
- [13] P. Paysan, R. Knothe, B. Amberg, S. Romdhani, and T. Vetter. A 3d face model for pose and illumination invariant face recognition. In *Proceedings of the 6th IEEE International Conference on Advanced Video and Signal based Surveillance (AVSS)*, pages 296–301. IEEE, 2009. [2](#)
- [14] E. Richardson, M. Sela, and R. Kimmel. 3d face reconstruction by learning from synthetic data. In *3D Vision (3DV), 2016 Fourth International Conference on*, pages 460–469. IEEE, 2016. [2](#)
- [15] T. Ritschel, T. Grosch, and H.-P. Seidel. Approximating dynamic global illumination in image space. In *Proceedings of the 2009 symposium on Interactive 3D graphics and games*, pages 75–82. ACM, 2009. [2](#)
- [16] S. Romdhani and T. Vetter. Efficient, robust and accurate fitting of a 3d morphable model. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 59–66. IEEE, 2003. [2](#)
- [17] S. D. Roth. Ray casting for modeling solids. *Computer graphics and image processing*, 18(2):109–144, 1982. [1](#)
- [18] S. Schönborn, B. Egger, A. Morel-Forster, and T. Vetter. Markov chain monte carlo for automated face image analysis. *International Journal of Computer Vision*, pages 1–24, 2016. [1](#), [2](#), [4](#), [6](#)
- [19] D. Shihlaei and V. Blanz. Realistic inverse lighting from a single 2d image of a face, taken under unknown and complex lighting. In *Automatic Face and Gesture Recognition (FG), 2015 11th IEEE International Conference and Workshops on*, volume 1, pages 1–8. IEEE, 2015. [2](#)
- [20] P.-P. Sloan, J. Kautz, and J. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 527–536. ACM, 2002. [1](#), [2](#), [5](#)
- [21] S. Suwajanakorn, I. Kemelmacher-Shlizerman, and S. M. Seitz. Total moving face reconstruction. In *European Conference on Computer Vision*, pages 796–812. Springer, 2014. [2](#)
- [22] S. Zhukov, A. Iones, and G. Kronin. An ambient light illumination model. In *Rendering Techniques 98*, pages 45–55. Springer, 1998. [2](#)