Chapter 19
3D Face Reconstruction from Single Two-Tone and Color Images

Ira Kemelmacher-Shlizerman, Ronen Basri, and Boaz Nadler

19.1 Introduction

This chapter addresses the problem of reconstructing the three-dimensional shape of faces from single images. We present an algorithm that uses prior knowledge of faces—a single shape model of a face—to eliminate the ambiguities in the reconstruction [13, 14]. The algorithm achieves veridical reconstruction results on images taken under a wide range of viewing conditions. In addition, it can reconstruct the shape of a face from two-tone ("Mooney") images of faces [15]. Our algorithm demonstrates the importance of "top-down" information in 3D shape reconstruction.

The extent to which internal representations affect perception is fundamental to the understanding of cognitive processes. Perceiving the appearance of a 3D shape in an image can be complicated as shapes are distorted by projection and their appearance is affected by lighting as well as by their color and texture. Yet people can readily perceive shape (perhaps qualitatively) merely from one image. A fundamental question therefore is whether the perception of shape is guided primarily by bottom-up processes, in which only image intensities are used along with generic assumptions regarding the statistics of natural scenes, or, alternatively, if it is dictated by top-down processes, which may be driven by memory and attention and

I. Kemelmacher-Shlizerman (✉)
Department of Computer Science and Engineering, University of Washington, Seattle, WA, USA
e-mail: kemelmi@cs.washington.edu

R. Basri · B. Nadler
Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel

R. Basri
e-mail: ronen.basri@weizmann.ac.il

B. Nadler
e-mail: boaz.nadler@weizmann.ac.il

S.J. Dickinson, Z. Pizlo (eds.), *Shape Perception in Human and Computer Vision*,
Advances in Computer Vision and Pattern Recognition,
DOI 10.1007/978-1-4471-5195-1_19,

Fig. 19.1 Two-tone ("Mooney") face images [1]. These images may initially seem difficult to interpret due to poor visual detail, but eventually lead to a rich and stable percept of a face

preceded by a preliminary recognition process. The example of random dot stereograms [12] suggests that the perception of 3D shape in stereo vision is governed by bottom-up processes. In contrast, two-tone images (see Fig. 19.1) suggest that familiarity with an object can enhance the perception of its shape.

Two-tone images were introduced in the 1950s by Craig Mooney [20] to investigate the development of shape perception in children [20, 28]. A number of recent studies suggest that the perception of Mooney images is driven by memory and attention and preceded by a preliminary recognition process. In particular, it was shown that people usually fail to perceive upside-down faces in two-tone images, arguably due to their unfamiliarity [8, 23], and that pre-exposure to original gray level (or color) image facilitates their recognition [6, 9]. Moreover, Moore and Cavanagh [21] showed that shape primitives (e.g., generalized cones) are rarely perceivable in two-tone images, both in isolation and in novel configuration with other primitives, even when the image contains explicit hints about the direction of the light source. These shapes, however, can readily be interpreted from gray level images and even from degraded line drawings. Familiar classes of objects, in contrast, are much more often perceivable in two-tone images. Even volumetric primitives of faces, if rearranged, cease to be perceived as coherent 3D objects. These findings support the view that the interpretation of Mooney images is guided by top-down processes. Here we provide further support for this claim by showing from a mathematical standpoint that, in the absence of a model, the interpretation of Mooney images is highly ambiguous.

Our further aim in this chapter is to provide an example of how top-down processing can play a role in the reconstruction of 3D faces. We focus on faces, as the overall similarity of faces [11] can provide a strong prior for reconstruction. Yet despite this similarity the reconstruction task is hard, since people are sensitive to minute shape differences across different individuals. 3D reconstruction of faces from single images can potentially be achieved by applying shape-from-shading (SFS) algorithms [10]. However, SFS requires knowledge of the lighting, albedo, and boundary conditions and is ill-posed in the absence of this information. The approach presented in this chapter uses prior knowledge about faces to achieve a well-posed formulation in which this missing information can be inferred. For a prior, we use a single 3D reference face model of either a different individual or a generic face. Our method assumes Lambertian reflectance (indeed, most face reconstruction methods assume that faces can be modeled accurately as Lambertian [19]), light sources at infinity, and rough alignment between the input image and the reference model. To model reflectance, we use a spherical harmonic approximation (following [3, 22]), which allows for multiple unknown light sources and attached

shadows. We note that other work used face priors (although not in the context of Mooney images), by combining information from hundreds of faces (e.g., [4, 5]). We find it interesting that veridical reconstructions can be obtained with just one model.

19.2 Reconstruction Ambiguities in Two-Tone Images

In this chapter, we ask whether a unique 3D shape can be recovered from a single Mooney image. We examine this question under typical SFS settings. We assume a single point light source of known magnitude and direction, the observed surface is Lambertian, and albedo is uniform (or otherwise known). Our formulation also accounts for boundary conditions. We focus on the *Mooney transition curve*, i.e., the boundary between bright and dark regions in the image, as the information contained in a Mooney image is captured almost entirely in this curve. We show that reconstruction is not unique even along this curve, indicating that top-down information is essential for shape perception.

Consider a gray level image $I(x, y)$ of a smooth Lambertian surface $z(x, y)$ with uniform albedo obtained with a directional illuminant $\mathbf{l} \in \Re^3$. The intensities $I(x, y)$ is given by $I = \mathbf{l}^T \mathbf{n}$, where $\mathbf{n} = n(x, y)$ denotes the surface normal at (x, y), $\mathbf{n} = (1/\sqrt{z_x^2 + z_y^2 + 1})(-z_x, -z_y, 1)$. A two-tone image is obtained from I by thresholding the image $I \geq T$ by some constant $T > 0$. Without loss of generality, we assume below that T is known, and that the light source direction coincides with the viewing direction, so that $\mathbf{l} = (0, 0, 1)$. Note however that our analysis can be applied to any directional source by a change of coordinates, as in [17], and the magnitude of the light can be scaled by appropriately scaling T. With these assumptions, we obtain

$$I(x, y) = \frac{1}{\sqrt{z_x^2 + z_y^2 + 1}}, \tag{19.1}$$

which can be expressed in the form of an Eikonal equation

$$|\nabla z|^2 = E(x, y) \tag{19.2}$$

on some closed domain $\Omega \subset \Re^2$, $E = (1/I^2) - 1$. Such an Eikonal equation can be solved for example by applying an upwind update scheme using a Dijkstra-like algorithm [17, 24, 27]. In general, such solutions require Dirichlet boundary conditions so that z needs to be specified at every local minimum of E (maximum of I) in Ω. These may include minimal points in Ω, as well as points along the boundaries of Ω. Our analysis therefore considers the introduction of boundary conditions.

Consider two surfaces z and z' that respectively produce two images I and I' (and hence E and E') which are "Mooney equivalent." By this we mean that $|\nabla z|^2 = E(x, y)$ and $|\nabla z'|^2 = E'(x, y)$ and $I = I' = \text{const}$ along an isoluminance curve γ. Some boundary conditions may also be specified, so that $z = z'$ (and at internal

points also $|\nabla z|^2 = |\nabla z'|^2 = 0$) in a set $\mathcal{B} \subset \Omega$. Let $\alpha(x, y) = z' - z$, our goal is given z to characterize the possible assignments of α along γ.

Subtracting the two eikonal equations for z and z' results in a new eikonal equation in α

$$|\nabla\alpha|^2 + 2\nabla\alpha \cdot \nabla z = E' - E. \tag{19.3}$$

To solve for α, we introduce a (local) change of coordinates $(x, y) \to (t, s)$ such that $\alpha_s = 0$ and $\alpha_t \neq 0$. In this coordinate frame, t points in the direction of the gradient of α, which is also the characteristic direction of (19.3). Consequently, (19.3) becomes

$$\alpha_t^2 + 2z_t\alpha_t - \left(E' - E\right) = 0. \tag{19.4}$$

This equation is quadratic in α_t and can have up to two real solutions,

$$\alpha_t = -z_t \pm \sqrt{z_t^2 + E' - E}. \tag{19.5}$$

We can use this equation to derive a general solution for α in the entire domain Ω by integrating (19.5) with respect to t along the characteristic directions as follows

$$\alpha(t, s) = -z(t, s) + z(t_0, s) \pm \int_{t_0}^{t} \sqrt{z_t^2 + E' - E}\, dt, \tag{19.6}$$

where the point $(t_0, s) \in \mathcal{B}$. One can readily verify that indeed $\alpha(t_0, s) = 0$.

Clearly, given a Mooney image we cannot use (19.6) to recover α, since in general neither E nor E' are known. However, along the transition curve, γ, we know that $E = E'$, and so (19.5) implies that $\alpha_t|_\gamma = -z_t \pm z_t \in \{0, -2z_t\}$. In general, we are interested here in the negative solution $\alpha_t|_\gamma = -2z_t$ since only this solution can produce a nontrivial ambiguity. We next use this solution to derive an explicit solution for α along the transition curve γ. Denote the arclength parameterization of γ by σ, and the angle between the tangent to γ and the t direction by $\theta(\sigma)$. Then

$$\alpha|_\gamma = -2\int_\gamma z_t \cos\theta\, d\sigma + \alpha(\sigma_0). \tag{19.7}$$

This solution implies that if we choose a set of characteristic directions for α along the transition curve γ then there can be exactly two shapes along this curve that are consistent with the input two-tone image, namely z and $z' = z + \alpha$. In general, however, we are free to choose any set of smoothly varying characteristic directions along γ and this way produce many additional solutions. A valid solution for α therefore must be consistent with the boundary conditions in $\mathcal{B}$, if such conditions are provided, and its gradients must coincide with *some* smoothly varying directional derivatives at points along γ. This implies in general that many ambiguities exist even if we only restrict our attention to the Mooney transition curve.

19.3 Shape Reconstruction with a Prior Model

In the remainder of this chapter, we introduce an algorithm for reconstructing the 3D shape of a face from a single image by exploiting our familiarity with faces as a class. Previous methods attempted to learn the set of allowable reconstructions from a large number of 3D laser-scanned faces. This was achieved by embedding all 3D faces in a linear space [2, 4, 26, 30] or by using a training set to determine a density function for faces [25, 29]. Similarly, Active Shape Models [5, 7, 18] seek to construct image-based, linear 2D representations of faces by exploiting large datasets of prototype faces for face recognition and image coding. In contrast to this work our method uses only a *single* reference model, and by that avoids the need to establish pointwise correspondence between many face models in a database.

Consider an image $I(x, y)$ of a face whose shape $z(x, y)$ is defined on a compact domain $\Omega \subset \Re^2$. We assume that the face is Lambertian with albedo $\rho(x, y)$, and that lighting can be an arbitrary combination of point sources, extended sources and diffuse lighting that need not be known ahead of time. Under these assumptions, Lambertian surfaces reflect only the low frequencies of lighting [3, 22], and so the reflectance function can be expressed in terms of spherical harmonics as

$$R\big(\mathbf{n}(x, y); \mathbf{l}\big) \approx \sum_{n=0}^{N} \sum_{m=-n}^{n} l_{nm} \alpha_n Y_{nm}\big(\mathbf{n}(x, y)\big), \tag{19.8}$$

where l_{nm} are the coefficients of the harmonic expansion of the lighting, α_n are factors that depend only on n and capture the effect of the Lambertian kernel acting as a low pass filter, so α_n becomes very small for large values of N, and $Y_{nm}(x, y)$ are the surface spherical harmonic functions evaluated at the surface normal. Because the reflectance of Lambertian objects under arbitrary lighting is in general very smooth, this approximation is highly accurate already when a low order (first or second) harmonic approximation is used.

For simplicity, we model the reflectance function using a first order harmonic approximation. In [14] we present a more general formulation using also the second order harmonics. We write the reflectance function in vector notation as

$$R\big(\mathbf{n}(x, y); \mathbf{l}\big) \approx \mathbf{l}^T \mathbf{Y}\big(\mathbf{n}(x, y)\big), \tag{19.9}$$

with $\mathbf{Y}(\mathbf{n}) = (1, n_x, n_y, n_z)^T$, where n_x, n_y, n_z are the components of the surface normal $\mathbf{n}$ and $\mathbf{l}$ is a four vector. The image irradiance equation is then expressed as $I(x, y) = \rho(x, y) R(x, y)$.

We are further given a reference face model and denote respectively by $z_{\text{ref}}(x, y)$, $\mathbf{n}_{\text{ref}}(x, y)$, and $\rho_{\text{ref}}(x, y)$ the surface, the normals, and the albedo of the reference face. We use the reference model to regularize the reconstruction problem. To that end, we define the difference shape and albedo as $d_z(x, y) = z(x, y) - z_{\text{ref}}(x, y)$ and $d_\rho(x, y) = \rho(x, y) - \rho_{\text{ref}}(x, y)$ respectively and require these differences to be smooth. We are now ready to define our optimization function:

$$\min_{\mathbf{l}, \rho, z} \int_{\Omega} \big(I - \rho \mathbf{l}^T \mathbf{Y}(\mathbf{n})\big)^2 + \lambda_1 (\triangle G * d_z)^2 + \lambda_2 (\triangle G * d_\rho)^2 \, dx \, dy, \tag{19.10}$$

where $\triangle G*$ denotes convolution with the Laplacian of a Gaussian, and λ_1 and λ_2 are positive constants. Below, we refer to the first term in this integral as the "data term" and the other two terms as the "regularization terms". Evidently, without regularization the optimization functional (19.10) is ill-posed. Specifically, for every choice of depth $z(x, y)$ and lighting vector $\mathbf{l}$ it is possible to prescribe albedo $\rho(x, y)$ to make the data term vanish. With regularization and appropriate boundary conditions, the problem becomes well-posed. Note that we chose to regularize d_z and d_ρ rather than z and ρ in order to preserve the discontinuities in z_{ref} and ρ_{ref}.

We assume the input image is roughly aligned to the reference model and approach this optimization by solving for lighting, depth, and albedo separately.

Step 1: Recovery of Lighting Coefficients In the first step, we attempt to recover the lighting coefficients $\mathbf{l}$, by fitting the reference model to the image. To this end, we substitute in (19.10) $\rho \to \rho_{\text{ref}}$ and $z \to z_{\text{ref}}$ (and consequently $\mathbf{n} \to \mathbf{n}_{\text{ref}}$). At this stage both regularization terms vanish, and only the data term remains:

$$\min_{\mathbf{l}} \int_{\Omega} \left(I - \rho_{\text{ref}} \mathbf{l}^T \mathbf{Y}(\mathbf{n}_{\text{ref}})\right)^2 dx\, dy. \tag{19.11}$$

In discrete form this produces a highly over-constrained linear least squares optimization system with only four unknowns, the components of $\mathbf{l}$, and can be solved simply using the pseudo-inverse. Our experiments indicate that, in practice, the error of recovering lighting using the face of a different individual is sufficiently small (around 4–6°).

Step 2: Depth Recovery We continue using $\rho_{\text{ref}}(x, y)$ for the albedo and turn to recovering $z(x, y)$. Below we exploit the reference face to further simplify the data term. The data term thus minimizes the squared difference between the two sides of the following system of equations

$$I = \rho_{\text{ref}} l_0 + \frac{\rho_{\text{ref}}}{N_{\text{ref}}}(l_1 z_x + l_2 z_y - l_3), \tag{19.12}$$

where $\mathbf{Y}(\mathbf{n}) = (1, z_x/N, z_y/N, -1/N)^T$ and we use $N_{\text{ref}}(x, y)$ to approximate $N(x, y) = \sqrt{z_x^2 + z_y^2 + 1}$. Replacing z_x and z_y, for example, by forward differences, the data term thus provides one equation for every unknown $z(x, y)$ (except for the pixels on the boundary of Ω). Note that by solving directly for $z(x, y)$ we in fact enforce consistency of the surface normals ("integrability"). Clearly, (19.12) is linear in $z(x, y)$ and so it can be solved using linear least squares optimization.

Next, we consider the regularization term $\lambda_1 \triangle G * d_z$. We implement this term as the difference between $d_z(x, y)$ and the average of d_z around (x, y) obtained by applying a Gaussian function to d_z. Consequently, this term minimizes the difference between the two sides of the following system of equations

$$\lambda_1\big(z(x, y) - G * z(x, y)\big) = \lambda_1\big(z_{\text{ref}}(x, y) - G * z_{\text{ref}}(x, y)\big). \tag{19.13}$$

This system too is linear in $z(x, y)$. The second regularization term vanishes since we have substituted ρ_{ref} for ρ.

Boundary Conditions for Depth Recovery For boundary conditions, we assume in our algorithm that the gradient of the surface in the direction perpendicular to the exterior boundary vanishes (i.e., the surface is planar near the boundaries; note that this does not imply that the entire bounding contour is planar). Specifically, we add for each boundary point the following constraint

$$\nabla z(x, y) \cdot \mathbf{n_c}(x, y) = 0. \tag{19.14}$$

where $\mathbf{n_c}(x, y)$ is a two-dimensional vector representing the normal to the bounding contour. These constraints will be roughly satisfied if the boundaries are placed in slowly changing parts of the face. They will be satisfied for example when the boundaries are placed along the cheeks and the forehead, but will not be satisfied when the boundaries are placed along the eyebrows, where the surface orientation changes rapidly.

Finally, since the obtained equation system involves only partial derivatives of $z(x, y)$, while $z(x, y)$ itself is absent from these equations, the solution can be obtained only up to an additive factor. We remedy this by arbitrarily setting one point to $z(x_0, y_0) = z_{\text{ref}}(x_0, y_0)$.

Step 3: Estimating Albedo Using the data term the albedo $\rho(x, y)$ is found by solving the following equation

$$I(x, y) = \rho(x, y)\mathbf{l}^T\mathbf{Y}(\mathbf{n}). \tag{19.15}$$

The first regularization term in (19.10) is independent of ρ, and so it can be ignored. The second term optimizes the following set of equations

$$\lambda_2 \triangle G * \rho = \lambda_2 \triangle G * \rho_{\text{ref}}. \tag{19.16}$$

These provide a linear set of equations, in which the first set determines the albedo values, and the second set smoothes these values. We avoid the need for boundary conditions simply by terminating the smoothing process at the boundaries.

19.4 Experiments

We demonstrate our algorithm on photographs taken under uncontrolled viewing conditions. Additional experiments and quantitative comparisons can be found in [14, 15]. We align the input image to the reference model by manually marking five points (the two centers of the eyes, the tip of the nose, the center of the mouth and the bottom of the chin), and then determine a 2×4 affine transformation, which aligns 3D points on the reference model to marked 2D points in the input image. After the alignment procedure the images are of typical size of 360×480 pixels. Our MATLAB implementation of the algorithm takes about 9 seconds on a quad-core AMD processor 2354 1100 MHz Linux workstation.

Fig. 19.2 A face model from the USF dataset used as a reference model in our experiments. The model is shown with uniform texture (*left*) and with an image overlay on the model (*right*)

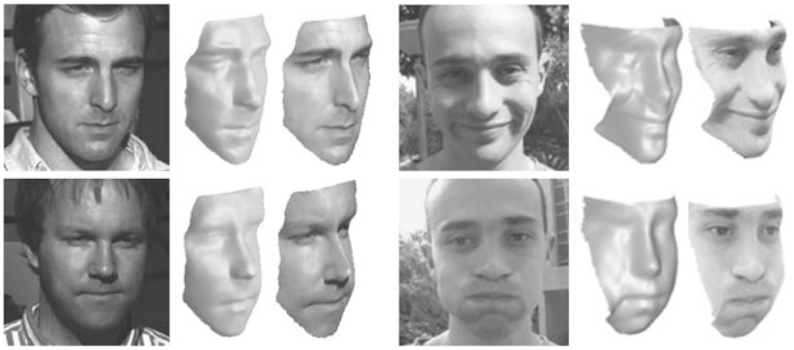

Fig. 19.3 Reconstruction results on images from the YaleB dataset (*left column*) and images photographed by us (*right column*). In each example we present the input image, our 3D shape reconstruction, and an image overlay on the reconstructed shape

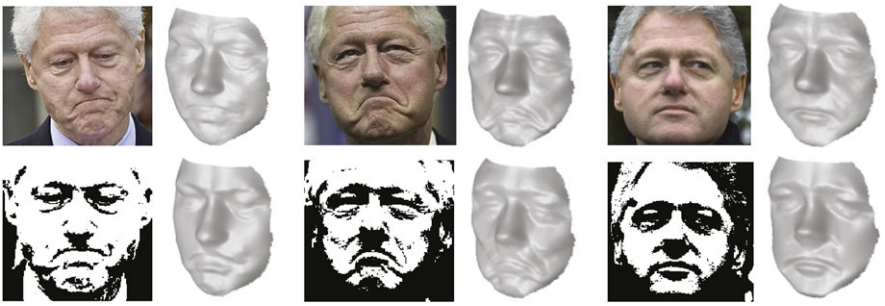

Fig. 19.4 Reconstruction from color (*top*) and two-tone (*bottom*) images. Each pair shows an input image and a reconstruction result

In Figs. 19.2–19.4, we show a few illustrative results obtained with our algorithm. The reference model used in our experiments is shown in Fig. 19.2. Figure 19.3 shows results on two images from the YaleB dataset and two more that were photographed by us. Figure 19.4 shows reconstruction results from three color images and from three two-tone images obtained by thresholding the intensity values.

We can see that convincing reconstructions are obtained for these images despite differences in identity, head orientation and facial expressions relative to the reference model. In addition, our reconstruction results for the two-tone images are similar to those obtained for the color images. These results are encouraging given the ill-posedness of the single view reconstruction problem and particularly reconstruction from two-tone images.

19.5 Conclusion

In this chapter, we explored the role of top-down information in the 3D reconstruction of faces from single images. We provided mathematical evidence that two-tone images provide ambiguous shape information and presented a novel method for reconstruction of faces from single image by using only a single reference model. Our results demonstrate that familiarity with faces as a class can help overcoming the difficulties in applying SFS algorithms and achieve veridical reconstructions even for Mooney images in which only two-tone intensity information is available.

We can foresee a number of potential directions to further extend over our method. One natural extension is to incorporate information from several images of the same individual, as in [16]; this could also address degeneracies of the current approach occurring under particular lighting conditions (see more details in [14]). Additionally, while facial expression is captured in the reconstruction, this is not directly targeted by the method and pixel-wise correspondence is not established. Establishing such correspondence can be useful also if we wish to generalize our ideas and apply them to other classes of objects. Finally, it will be interesting to further explore to what extent our algorithmic objectives are indeed achieved by the visual cortex.

Acknowledgement Research was supported in part by the Israel Science Foundation grant number 266/02 and by the European Commission Project IST-2002-506766 Aim Shape. The vision group at the Weizmann Institute is supported in part by the Moross Laboratory for Vision Research and Robotics.

References

1. http://www.princeton.edu/artofscience/gallery/view.php%3fid=77.html
2. Atick JJ, Griffin PA, Redlich AN (1996) Statistical approach to shape from shading: reconstruction of 3d face surfaces from single 2d images. Neural Comput 8(6):1321–1340
3. Basri R, Jacobs DW (2003) Lambertian reflectance and linear subspaces. PAMI 25(2)
4. Blanz V, Vetter TA (1999) A morphable model for the synthesis of 3d faces. Comput Graph I:187–194
5. Cootes TF, Edwards GJ, Taylor CJ (1998) Active appearance models. In: ECCV, vol 2
6. Dolan RJ, Fink GR, Rolls E, Booth M, Holmes A, Frackowiak RSJ, Friston KJ (1997) How the brain learns to see objects and faces in an impoverished context. Nature 389
7. Edwards GJ, Lanitis A, Taylor CJ, Cootes TF (1996) Modelling the variability in face images. In: Proc of the 2nd int conf on automatic face and gesture recognition, vol 2

8. George N, Jemel B, Fiori N, Renault B (1997) Face and shape repetition effects in humans: a spatio-temporal ERP study. NeuroReport 8(6):1417–1423
9. Hegde J, Thompson S, Kersten D (2007) Identifying faces in two-tone ('mooney') images: a psychophysical and fMRI study. J Vis 7(9):624
10. Horn BKP, Brooks MJ (eds) (1989) Shape from shading. MIT Press, Cambridge
11. Hursh TM (1976) The study of cranial form: measurement techniques and analytical methods. In: Giles E, Fiedlaender J (eds) The measures of man
12. Julesz B (1971) Foundations of cyclopean perception. The University of Chicago Press, Chicago
13. Kemelmacher I, Basri R (2006) Molding face shapes by example. In: ECCV, vol 1, pp 277–288
14. Kemelmacher I, Basri R (2011) 3d face reconstruction from a single image using a single reference face shape. IEEE Trans Pattern Anal Mach Intell 33(2):394–405
15. Kemelmacher-Shlizerman I, Basri R, Nadler B (2008) 3d shape reconstruction of mooney faces. In: CVPR
16. Kemelmacher-Shlizerman I, Seitz SM (2011) Face reconstruction in the wild. In: ICCV
17. Kimmel R, Sethian JA (2001) Optimal algorithm for SFS and path planning. J Math Imaging Vis 14(3):237–244
18. Lanitis A, Taylor CJ, Cootes TF (1997) Automatic interpretation and coding of face images using flexible models. IEEE Trans Pattern Anal Mach Intell 19(7):743–756
19. Marschner SR, Westin SH, Lafortune EPF, Torrance KE, Greenberg DP (1999) Image-based brdf measurement including human skin. In: 10th eurographics workshop on rendering, pp 139–152
20. Mooney CM (1957) Age in the development of closure ability in children. Can J Psychol 11:219–226
21. Moore C, Cavanagh P (1998) Recovery of 3d volume from 2-tone images of novel objects. Cognition 67:45–71
22. Ramamoorthi R, Hanrahan P (2001) On the relationship between radiance and irradiance: Determining the illumination from images of a convex Lambertian object. J Opt Soc Am 18(10):2448–2459
23. Rodriguez E, George N, Lachaux JP, Martinerie J, Renault B, Varela FJ (1999) Perception's shadow: long-distance synchronization of human brain activity. Nature 397:430–433
24. Sethian JA (1996) A fast marching level set method for monotonically advancing fronts. Proc Natl Acad Sci USA 93:1591–1595
25. Sim T, Kanade T (2001) Combining models and exemplars for face recognition: an illuminating example. In: CVPR workshop on models versus exemplars
26. Smith WAP, Hancock ER (2005) Recovering facial shape and albedo using a statistical model of surface normal direction. In: ICCV
27. Tsitsiklis JN (1994) Efficient algorithms for globally optimal trajectories. In: Proc conf on decision and control
28. Yoon J, Winawer J, Wittoft N, Markman E (2007) Mooney image perception in preschool-aged children. J Vis 7(9):548
29. Zhang L, Samaras D (2003) Face recognition under variable lighting using harmonic image exemplars. In: CVPR
30. Zhou SK, Chellappa R, Jacobs DW (2004) Characterization of human faces under illumination variations using rank, integrability, and symmetry constraints. In: ECCV