

ARTINT 1001

Comment on “Numerical shape from shading and occluding boundaries”

K. Ikeuchi

School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA

Abstract

Ikeuchi, K.. Comment on “Numerical shape from shading and occluding boundaries”, *Artificial Intelligence* 59 (1993) 89–94.

The paper “Numerical shape from shading and occluding boundaries” [10] proposed a method to determine the 3D shape of an object from a single brightness image using the smoothness constraint on surface orientations, and the orientations of the surface given along an occluding boundary. This paper contributed both a specific algorithm for solving a specific problem, and a general approach that could be applied in other low-level vision problems.

1. Research background

It has long been known that the human visual system is capable of determining the 3D structure (relative depth or surface orientations) of a scene, given only 2D information. For example, a black-and-white photograph of an object is only 2D, yet can be used by a human to determine the 3D shape of the object. The human visual system uses a variety of cues to determine 3D shape. One such cue is shading information: the way in which brightness varies along the surface of an object is used by the human visual system to determine the underlying 3D shape of the object; the process is known as shape-from-shading.

Correspondence to: K. Ikeuchi, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA.

In the late 1970s and early 1980s, at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology, several projects were conducted to investigate computational algorithms to emulate on computers these modules of the human visual system that determine (relative) depth from images. Such projects investigated shape-from-shading (Horn, Woodham, Ikeuchi, Pentland), shape-from-texture (Richard, Witkin, Stevens), optical flow (Horn, Marr, Schunk, Hildreth), and binocular stereo (Marr, Poggio, Grimson).

Shape-from-shading research, originally suggested as a thesis topic by Minsky, was initiated by Horn as his Ph.D. thesis in 1970 [7]. Horn formalized the shape-from-shading problem using the nonlinear partial differential *image irradiance equation*,

$$E(x, y) = R(p, q), \quad (1)$$

where $R(p, q)$ is a known reflectance function of an object surface, represented using the standard (p, q) gradient space coordinate system to define unknown orientations and assuming orthographic projection down the z axis onto the (x, y) plane for image formation, and $E(x, y)$ is a given input brightness distribution at each image point, (x, y) .

Horn solved the equation using the characteristic strip expansion method. The method draws characteristic strips—special strips which satisfy a certain condition—in a shaded image and determines orientations iteratively along the characteristic strips from a point whose orientation is known. As the starting point, Horn used the point which provides the maximum brightness, $R(p, q) = 1$. (In particular, on a Lambertian surface, the point of the maximum brightness is oriented toward the light source.) The method demonstrated was able to determine the shape of 3D objects such as a human nose from its black/white shaded image.

Horn's original method, however, has two main defects.

- Direct integration of the characteristic equations to find the characteristic strips suffers from noise sensitivity in practical implementations. As the solution proceeds along constructed strips from the initial points, these strips, as well as the computed orientations, can deviate from the actual strip positions and orientations as a result of quantization errors and other noise influences.
- The method starts characteristic strips from a point with a known orientation. The method employs the point of maximum brightness for the starting point. However, at such points the partial derivatives of R vanish, so the strips cannot be uniquely constructed. To overcome this problem, the original method uses the points on a circumscribing circle to approximate the starting points.

2. Technical contents

In 1979¹ we (Ikeuchi and Horn) solved these problems by introducing the *smoothness constraint* and by utilizing the information given along *occluding boundaries*:

- *Noise sensitivity*: Horn's original method can only determine surface orientations along the characteristic strips. Thus, a single noisy point on a strip could throw off all of the following computations. Our method introduced the smoothness constraint, which required neighboring points to have similar surface orientations:

$$p_x^2 + p_y^2 + q_x^2 + q_y^2 = 0. \quad (2)$$

By using this constraint, our method propagated, combined, and averaged orientations over neighboring grid points to determine the orientation at the central point. By averaging the information in a neighborhood, sensitivity to noise was greatly reduced.

- *Singular point*: Horn's original method can only use the orientation at a singular point. Our method uses more widely available orientations along an occluding boundary. (Horn's method cannot use these orientations due to its parameterization.)

At the points along occluding boundaries of objects, orientations are perpendicular to the lines of sight. From the shape of the silhouette, which is the projection of an occluding boundary onto the image plane, we can determine the tangential direction of the occluding boundary parallel to the image plane, and thus, the orientation there. Our method uses these orientations as the initial conditions for the algorithm.

The image irradiance equation, $E(x, y) - R(p, q) = 0$, only provides one constraint at each image point. The surface orientation of the point has two degrees of freedom (p, q) . We used the smoothness constraint, $p_x^2 + p_y^2 + q_x^2 + q_y^2 = 0$, as an additional constraint. (Horn's method uses a weak smoothness constraint given by the continuity of the characteristics strips.)

We formalized the shape-from-shading problem as a minimization problem of the difference between the observed brightness and the expected brightness given through the image irradiance equation from the expected surface orientation, while maintaining the smoothness constraint.

$$\int \int_A (E(x, y) - R(p, q))^2 + \lambda(p_x^2 + p_y^2 + q_x^2 + q_y^2) \, dx \, dy. \quad (3)$$

¹The original working memo was published November 1979, the revised AI memo published February 1980, and the revised journal article was published in August 1981.

The minimization is performed over the region, A , enclosed by an occluding boundary, ∂A ; the aim is to obtain a distribution of (p, q) s which minimize the equations over the region, A .

Using the calculus of variations, we formulated an iterative scheme to obtain the (p, q) which minimizes the equation. Since the iterative scheme requires a boundary condition, we used the orientations along the occluding boundary, ∂A . The method was demonstrated to determine orientations over the region, A , from the shading information over the region, A , and the orientation along the boundary, ∂A .

3. Contributions

One of the main intellectual contributions of the paper came from reporting the effectiveness of the smoothness constraint on low-level vision modules. The smoothness constraint forces neighboring points to have similar surface orientations. Later this constraint was generalized into the *regularization term* by Poggio. The paper predicted that the smoothness constraint would be effective with other low-level vision modules. (See the discussion in [10, p. 181].)

The paper also demonstrated the effectiveness of the variational approach for low-level vision problems. Using this method, various low-level vision problems can be solved in a uniform manner.

4. Open issues

The following are research issues that were left out of the original paper but are still open. Some of the recent representative research toward the issues are also discussed:

The smoothness constraint

Though robust, the method based on the smoothness constraint does deviate from the real solution. That is, the algorithm converges to shapes that are slightly different from the real shapes. Another problem of the smoothness constraint is that it does not guarantee that the solution obtained is integrable because our method determines p and q independently. Several researchers (Horn and Brooks [5,6] and Frankot and Chelappa [4]) pointed out these two effects and proposed new algorithms by introducing integrability constraints. However, I (Ikeuchi) believe that the smoothness constraint is still interesting enough as a mechanism to emulate modules of the human visual system.

- (1) The smoothness constraint can be applied to a wide variety of low-level vision modules regardless of parameterizations. Some of the proposed constraints are only valid for a certain parameterization. The notion of smoothness should be independent of parameterizations. It is quite likely that the human visual system imposes such smoothness constraints (nearby points have similar characteristics) to its default interpretation of a scene.
- (2) The method based on the smoothness constraint has the preferred default shape. In the absence of any shading information, the method applied to the boundary condition yields the default shape. When shading information is present, this shape is modified so as to satisfy the image irradiance equation. It seems to me that this mechanism is quite reasonable as the behavior of the human visual system. Namely, the mechanism has an underlying default shape and outputs such a default shape when no input information is available. When some input is available, it modifies the output accordingly.

From these two points, I conjecture that some sort of smoothness constraint mechanism exists in the human visual system. Research on the similarities and differences between human and machine shape-from-shading algorithms is an interesting topic. Such comparisons are reported by Mingolla and Todd [11].

Occluding boundary

The method requires surface orientations along a closed occluding boundary. We can observe a whole occluding boundary only when a viewer and a light source are located at the same position. Under general conditions, an illuminated region is partly enclosed by an occluding boundary and partly by a self-shadowed boundary. Along a self-shadowed boundary, no orientation information is available. Thus, under general conditions, our method cannot be applied to the shape-from-shading problem. It is necessary to remedy this point (see Pentland [14]).

Binocular stereo provides depth information along a boundary, ∂A . In order to obtain depth information over the region A , it is necessary to interpolate depth information from that along the boundary ∂A . Shape-from-shading provides (relative) depth information over the region A , from those along the boundary ∂A . Both modules are complementary in nature, and thus combining the methods is an interesting research goal (Ikeuchi [8,9], Blake et al. [1]).

Known reflectance map

Our shape-from-shading algorithm requires a known exact reflectance map, that is, known reflectance characteristics, and a known light source direction.

It is rather unusual to know the exact reflectance characteristics of an object and the exact direction of the light source. A method for determining the shape using a rough reflectance map is necessary (Brooks and Horn [2], Nayar et al. [12]).

Interreflection

Finally, the method assumes that one light source directly illuminates a surface. However, it often occurs that several light sources illuminate a surface simultaneously, and interreflections occur among nearby surfaces. Research on shape-from-shading under complicated illumination conditions should be explored (Forsyth [3], Nayar et al. [13]).

References

- [1] A. Blake, A. Zisserman and G. Knowles, Surface descriptions from stereo and shading, *Image Vis. Comput.* 3 (4) (1985) 183–191.
- [2] M.J. Brooks and B.K.P. Horn, Shape and source from shading, in: *Proceedings IJCAI-85*, Los Angeles, CA (1985) 932–936.
- [3] D. Forsyth and A. Zisserman, Reflections on shading, *IEEE Trans. Pattern Anal. Mach. Intell.* 13 (7) (1991) 671–679.
- [4] R.T. Frankot and R. Chellappa, A method for enforcing integrability in shape from shading algorithm, *IEEE Trans. Pattern Anal. Mach. Intell.* 10 (4) (1988) 439–451.
- [5] B.K.P. Horn, Height and gradient from shading, *Int. J. Comput. Vis.* 5 (1) (1990) 37–74.
- [6] B.K.P. Horn and M.J. Brooks, The variational approach to shape from shading, *Comput. Vis. Graph. Image Process.* 33 (2) (1986) 174–208.
- [7] B.K.P. Horn, Shape from shading: a method for obtaining the shape of a smooth opaque object from one view, Tech. Rept. MAC-TR-79, MIT, Cambridge, MA (1970).
- [8] K. Ikeuchi, Reconstructing a depth map from intensity map, in: *Proceedings International Conference on Pattern Recognition* (1984) 736–738.
- [9] K. Ikeuchi, Determining a depth map using a dual photometric stereo, *Int. J. Rob. Res.* 6 (1) (1987) 15–31.
- [10] K. Ikeuchi and B.K.P. Horn, Numerical shape from shading and occluding boundaries, *Artif. Intell.* 17 (1981) 141–184.
- [11] E. Mingolla and J.T. Todd, Perception of solid shape from shading, *Biol. Cybern.* 53 (1986) 137–151.
- [12] S.K. Nayar, K. Ikeuchi and T. Kanade, Determining shape and reflectance of hybrid surfaces by photometric sampling, *IEEE Trans. Rob. Autom.* 6 (4) (1990) 418–431.
- [13] S.K. Nayar, K. Ikeuchi and T. Kanade, Shape from interreflections, *Int. J. Comput. Vis.* 6 (3) (1991) 173–195.
- [14] A.P. Pentland, Linear shape from shading, *Int. J. Comput. Vis.* 4 (1990) 153–162.