

Relating Image Warping to 3D Geometrical Deformations

A.L. Yuille¹, Mario Ferraro², and Tony Zhang³

¹ Smith-Kettlewell Eye Research Institute, San Francisco, CA 94115.

² Dipartimento di Fisica Sperimentale, Università' di Torino, via Giuria 1, 10125 Torino, Italy.

³ Division of Applied Sciences, Harvard University, Cambridge MA, 02138

Abstract. We demonstrate that, for a large class of reflectance functions, there is a direct relationship between image warps and the corresponding geometric deformations of the underlying three-dimensional objects. This helps explain the hidden geometrical assumptions in object recognition schemes which involve two-dimensional image warping computed by matching image intensity. In addition, it allows us to propose a novel variant of shape from shading which we call shape from image warping. The idea is that the three-dimensional shape of an object is estimated by determining how much the image of the object is warped with respect to the image of a known prototype shape. Therefore detecting the image warp relative to a prototype of known shape allows us to reconstruct the shape of the imaged object. We derive properties of these shape warps and illustrate the results by recovering the shapes of faces.

1 Introduction

Recent work on object recognition [6] uses two-dimensional warping of intensity images to allow for the changing three-dimensional geometry of the objects. For example, the image of a viewed object will change with the angle of view. For small changes of angle, or small deformations of shape, the change can be modeled as a spatial warp of the intensity image. The approach makes the implicit assumption that we can model geometrical changes in three-dimensions by warps in the two-dimensional image plane. When is this assumption valid?

In this paper we show that, for a large class of reflectance functions, there is a direct relationship between image warps and geometrical changes of the underlying three-dimensional objects. We also demonstrate that not all warps are physically reasonable and determine constraints that physical warps must satisfy.

Our analysis explains the hidden assumptions used by Hallinan and other workers on object recognition [6]. It allows us to understand the relationship to other object recognition theories based on three-dimensional geometry.

In addition, it offers a novel approach to shape from shading [7], [8], [11]. We call our method *shape from warping*. By contrast to standard techniques, our approach works by assuming prototype models of shape and by estimating the warp between the input image and the image of the prototype. From these image

warps we show how to deduce the three-dimensional shape. This method allows us to introduce object specific knowledge into the shape estimation. We will prove that this approach works without needing to know the precise reflectance function of the object.

Our approach has some similarities to recent work [1], [3], which shows that shape can be recovered from an intensity image provided prior assumptions are made about the shape class. The techniques used in this paper, however, are very different.

This paper is organized as follows. Section (2) demonstrates the basic relationship between two-dimensional image warping and three-dimensional geometrical variations. In Section (3) we describe how the surface integrability condition put constraints on the class of warps by requiring that they generate consistent surfaces. In Section (4) we discuss the underlying assumptions of object recognition theories which involve two-dimensional image warping. Section (5) illustrates shape from warping by recovering the shape of faces.

2 Relating Image and Geometrical Warps

Suppose we have a surface, or object, of the form $z(\mathbf{x}) = f(\mathbf{x};)$. (More precisely, we consider a *Monge patch*, which is assumed to be smooth, C^∞ , or at least C^2 .) This surface corresponds to an intensity image $I(\mathbf{x})$. The image is related to the surface by a reflectance function, see [7]. For example, it is common to assume that:

$$I(\mathbf{x}) = R(\mathbf{n}(\mathbf{x}), \mathbf{k}, \mathbf{s}), \quad (1)$$

where R is the reflectance function, $\mathbf{n}(\mathbf{x})$ are the surface normals, \mathbf{k} is the viewer direction, and \mathbf{s} is the light source direction. (This is the most general class of reflectance function in common use and includes, for example, the Phong model as a special case provided the albedo is constant).

Our key result is that *warping the image corresponds to warping the surface normals of the underlying surface*. More precisely, suppose we apply a warp $\phi(\cdot)$ to the image by the mapping:

$$\phi : \mathbf{x} \rightarrow \phi(\mathbf{x}), \quad (2)$$

then this induces a warp:

$$I(\mathbf{x}) \mapsto I(\phi(\mathbf{x})), \quad (3)$$

to the image and, by equation (1), a warp:

$$\mathbf{n}(\mathbf{x}) \mapsto \mathbf{n}(\phi(\mathbf{x})), \quad (4)$$

to the *surface normals* of the surface.

This result can be used in two ways. Firstly, it can be used to explain the hidden geometrical assumptions of theories of object recognition which make use of two dimensional image warps [6], see section (4) for more details. Secondly, the result can be exploited to develop a novel approach to shape from shading which we call

shape from warping. If the warping assumption is valid then we can use the image warps to recover three-dimensional shape assuming that we have a known three-dimensional prototype with surface normals $\{\mathbf{n}_0(\mathbf{x})\}$ and a corresponding image $I_0(\mathbf{x})$. For a given image $I(\mathbf{x})$ we find the warp $\phi(\mathbf{x})$ such that $I(\mathbf{x}) = I_0(\phi(\mathbf{x}))$ and hence determine the shape to be $\mathbf{n}(\mathbf{x}) = \mathbf{n}_0(\phi(\mathbf{x}))$.

It should be emphasized that the approach assumes that the reflectance function is of form $I(\mathbf{x}; \alpha) = R(\mathbf{n}(\mathbf{x}; \alpha), \mathbf{k}, \mathbf{s})$. This will not be true if the object has significant albedo changes but may be a reasonable approximation for many objects. (We note that almost all shape from shading algorithms assume constant albedo). This seems to be true for our experiments on faces, see section (5).

But what warps are allowable? From equation (4) we see that they warp the surface normals of a prototype shape, and not all such warps will form consistent surfaces. This is investigated in the next section.

3 Normal Consistency

The preceding section showed that image warps often corresponded to warps of surface normals. In what situations will this give a consistent surface? Warps transform a given set of surface normals \mathcal{N} , into a different set say \mathcal{N}' , but this does not guarantee the existence of a consistent surface with these normals. To ensure the existence of a such a surface a further condition must be satisfied, that is referred to as the integrability condition [12]. Let us derive the *surface integrability condition* (see, for example, [8]).

Suppose we have a surface $z = f(x, y)$. Then its surface normals are defined by:

$$\mathbf{n}(x, y) = \frac{1}{\{1 + \nabla f \cdot \nabla f\}^{1/2}}(-f_x, -f_y, 1), \quad (5)$$

where $\nabla f = (f_x, f_y)$. We can write $\mathbf{n} = (n_1, n_2, n_3)$ and observe that:

$$\frac{n_1}{n_3} = -f_x, \quad \frac{n_2}{n_3} = -f_y. \quad (6)$$

Hence we derive the *surface integrability condition*:

$$\frac{\partial}{\partial y} \left(\frac{n_1}{n_3} \right) = \frac{\partial}{\partial x} \left(\frac{n_2}{n_3} \right). \quad (7)$$

We have shown that this is a necessary condition for the surface to be consistent. To see that it is sufficient we observe that, by elementary vector calculus [4], equation (7) implies that there exists a function $\psi(x, y)$ such that:

$$\left(\frac{n_1}{n_3} \right) = \frac{\partial \psi}{\partial x}, \quad \left(\frac{n_2}{n_3} \right) = \frac{\partial \psi}{\partial y}. \quad (8)$$

We can solve equation (8) for n_1, n_2, n_3 , using the normalization condition $n_1^2 + n_2^2 + n_3^2 = 1$, and obtain the surface expression (5) after identifying ψ with $-f$. Thus we see that equation (7) is also a sufficient condition.

We must now tackle the harder task of putting consistency conditions on the warp so that the warped normals are consistent. Suppose $\mathbf{n}(\mathbf{x})$ are the surface normals of the prototype surface, and hence obey the surface integrability condition. Let us apply a warp $\phi(\mathbf{x}) = (\phi_1(\mathbf{x}), \phi_2(\mathbf{x}))$ to the prototype surface. The resulting surface $\mathbf{n}(\phi(\mathbf{x}))$ is consistent provided:

$$\frac{\partial}{\partial y} \left(\frac{n_1(\phi(\mathbf{x}))}{n_3(\phi(\mathbf{x}))} \right) = \frac{\partial}{\partial x} \left(\frac{n_2(\phi(\mathbf{x}))}{n_3(\phi(\mathbf{x}))} \right). \quad (9)$$

By equation (6) we have:

$$\begin{aligned} \frac{n_1(\phi(\mathbf{x}))}{n_3(\phi(\mathbf{x}))} &= -\frac{\partial f}{\partial x(\phi(\mathbf{x}))}, \\ \frac{n_2(\phi(\mathbf{x}))}{n_3(\phi(\mathbf{x}))} &= -\frac{\partial f}{\partial y(\phi(\mathbf{x}))}, \end{aligned} \quad (10)$$

where the derivatives on the right hand sides are evaluated at $\phi(\mathbf{x})$ (a example is given below for readers unfamiliar with this notation). We observe that

$$\frac{\partial}{\partial y} \left(\frac{\partial f}{\partial x(\phi(\mathbf{x}))} \right) = \frac{\partial^2 f}{\partial x^2(\phi(\mathbf{x}))} \frac{\partial \phi_1}{\partial y} + \frac{\partial^2 f}{\partial x \partial y(\phi(\mathbf{x}))} \frac{\partial \phi_2}{\partial y}. \quad (11)$$

A similar result holds if we replace x by y .

Substituting from (10) into (9) and using (11) yields the result:

$$\begin{aligned} \phi_{1,y} \left\{ \frac{\partial^2 f}{\partial x^2} \right\}_{\phi(\mathbf{x})} + \phi_{2,y} \left\{ \frac{\partial^2 f}{\partial x \partial y} \right\}_{\phi(\mathbf{x})} = \\ \phi_{1,x} \left\{ \frac{\partial^2 f}{\partial x \partial y} \right\}_{\phi(\mathbf{x})} + \phi_{2,x} \left\{ \frac{\partial^2 f}{\partial y^2} \right\}_{\phi(\mathbf{x})}. \end{aligned} \quad (12)$$

This gives a relationship between the Hessian of the surface (terms $\partial^2 f / \partial x^2$, etc.) and the *stress tensor* of the warps which is defined to have components $\phi_{1,x}, \phi_{1,y}, \phi_{2,x}, \phi_{2,y}$.

We can get some clarification of the notation used in this derivation by considering the one-dimensional case where the surface $f(x)$ can be expressed as a power series:

$$f(x) = \sum_{r=0}^N a_r x^r, \quad (13)$$

and warping it by $\phi(x)$ gives:

$$f(\phi(x)) = \sum_{r=0}^N a_r \{\phi(x)\}^r. \quad (14)$$

In this case we find that:

$$\frac{\partial f}{\partial x(\phi(x))} = \sum_{r=0}^N a_r r \{\phi(x)\}^{r-1} = \frac{\partial}{\partial \phi(x)} f(\phi(x)). \quad (15)$$

Some intuition can be obtained by considering surfaces of revolution. These surfaces are essentially one-dimensional and it can be shown [13] that warps which preserve the revolution property will satisfy this equation.

4 Underlying Geometric Assumptions of Image Warping

The Hallinan model for face recognition [6] assumes that the image of a given face can be written as⁴:

$$I(\mathbf{x}) = \sum_{i=1}^5 \alpha_i B_i(\phi(\mathbf{x})), \quad (16)$$

where the $\{B_i(\cdot)\}$ are lighting basis functions [5] – which model the appearance of the face under different lighting conditions – and $\phi(\cdot)$ is a geometric warp which is required to be monotonic. The $\{\alpha_i\}$ are coefficients corresponding to the specific lighting conditions.

The Hallinan model therefore assumes that any image of a face can be obtained by monotonically warping a prototype face model. Now the albedo is approximately constant over most of the face and, in particular, it is constant over those regions of the face where most of the shape changes occur – the cheeks, forehead, and nose. Thus we argue that faces approximately satisfy the assumptions of shape from warping.

Therefore the image warps of the Hallinan model should correspond directly to warps of the surface normals of faces⁵. If so, the monotonicity assumption of Hallinan’s warps can be shown to imply that no surface extrema can be created in the images objects [13]. In other words, the model given by equation (16) will break down if some of the objects have a different number of surface extrema than the prototype. This will not happen for faces but it will be a problem if the model is applied to other object classes.

It is interest to contrast Hallinan’s face model with the one recently proposed by Atick et al. [1]. This model represents the face in terms of a three-dimensional model which is generated by doing principal component analysis of a set of three-dimensional range data of faces. A face can therefore be represented by its coefficients of expansion in a basis of *eigenheads*. By contrast, Hallinan has computed a class of two-dimensional eigenwarps. Our paper therefore suggests that Hallinan’s image eigenwarps correspond to eigenwarps of the surface normals and hence are closely related to Atick’s eigenheads.

5 Examples of Shape from Warping

We illustrate our theory on real images of faces obtained from [6]. Our prototype model consists of a face image $I_p(\mathbf{x})$ and its corresponding surface shape $z = f_p(\mathbf{x})$ obtained from laser-range data. From this surface we compute the surface normals $\mathbf{n}_p(\mathbf{x})$.

For given input images $\{I_\alpha(\mathbf{x})\}$ we compute the spatial warps $\phi_\alpha(\mathbf{x})$ relative to the prototype image by using the algorithm, and code, described in [6], see

⁴ The full model also includes a global affine warp which we will neglect for simplicity.

⁵ This will be verified by the experiments described in the following section.

figure (1). These warps are calculated to minimize a matching energy function:

$$E[\phi] = \frac{1}{|D|} \int_D \psi\{I(\phi(\mathbf{x})) - I_p(\mathbf{x})\} dx + \int_D (\text{tr} J(\mathbf{x})^T J(\mathbf{x}) \{1 + \frac{1}{\det J(\mathbf{x})^2}\} - 4) dx, \quad (17)$$

where J is the Jacobian matrix of the warp field $\phi(\mathbf{x})$ and ψ is a robust norm [9].

More sophisticated versions of the algorithm in [6] will obtain the spatial warps even if the lighting conditions are unknown. But we will not deal with this case in this paper.

The code described in [6] does not put any restrictions on the class of warps other than requiring monotonicity. It would be interesting to impose the consistency condition, see equation (12), while calculating the warps but this is difficult because the warps appear in the arguments of the surface derivative terms.

Instead we use the algorithm to compute the image warps $\phi(\mathbf{x})$, calculate the warps normal fields $\mathbf{n}(\mathbf{x}) = \mathbf{n}_p(\phi(\mathbf{x}))$ and then determine the closest consistent surface by minimizing the cost function:

$$E[z] = \int \{(\frac{\partial z}{\partial x} - \frac{n_1(\mathbf{x})}{n_3(\mathbf{x})})^2 + (\frac{\partial z}{\partial y} - \frac{n_2(\mathbf{x})}{n_3(\mathbf{x})})^2\} dx. \quad (18)$$

The resulting warps and reconstructed faces are shown in figures (1, 2). The warps are generally well calculated by Hallinan's algorithm though close inspection, see [13], shows some errors, particularly on the right boundary of object number 3. These errors become apparent in the reconstructions, see figure (2), and we see that errors arise on the edge of object 3.



Fig. 1. Applying our algorithm to the first object. Left the intensity of the first object, center warping the prototype to match the object, right the estimated shape of the first object.

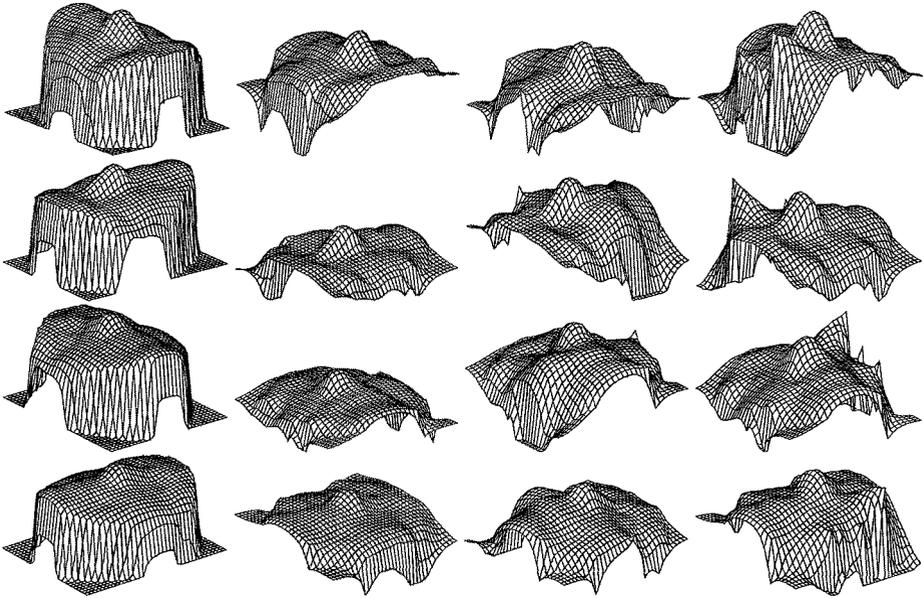


Fig. 2. Showing the estimated face for the objects, from different viewpoints, and comparing them to the prototype shape. The four rows correspond to different viewpoints. The first column shows the prototype shape and the next three columns show the three objects. Observe the errors which sometimes occur at the object boundaries. These are side effects from the warping algorithm.

It seems, therefore, that the errors in reconstruction are due to partial failures in the warping algorithm. Thus, since these are nontrivial images, we regard this as proof of concept for shape from warping.

6 Summary

This paper has demonstrated a fundamental relationship between two-dimensional image warps and three-dimensional geometrical transformations.

Our analysis shows the underlying geometrical assumptions about three-dimensional shape used in models of object recognition which rely on dense two-dimensional warps computed by matching image intensity [6].

We also described *shape from warping*, a variant of shape from shading, which

allows shape to be determined for a class of surfaces without needing to know the exact reflectance function. We investigated the class of warps that are consistent and what types of surfaces could be reconstructed in this way. The theory was illustrated by applying it to recovering the shape of faces.

Acknowledgements

We would like to thank Peter Belhumeur, Gaile Gordon, and particularly Peter Hallinan for access to their computer code. Support was provided by NSF Grant IRI 93-17670 and ARPA/ONR Contract N00014-95-1-1022.

References

1. J.J. Atick, P.A. Griffin, and A.N. Redlich. "Statistical Approach to Shape from Shading: Reconstruction of 3D Face Surfaces from Single 2D Images". Preprint. The Computational Neuroscience Laboratory. The Rockefeller University. New York, NY. 1995.
2. R. Epstein, P.W. Hallinan and A.L. Yuille. " $5 \pm$ Eigenimages Suffice: An Empirical Investigation of Low-Dimensional Lighting Models". In *Proceedings of IEEE WORKSHOP ON PHYSICS-BASED MODELING IN COMPUTER VISION*. 1995.
3. R. Epstein, A.L. Yuille, and P.N. Belhumeur. "Learning object representations from lighting variations". In **Object Representation in Computer Vision II**. Eds. J. Ponce, A. Zisserman, and M. Hebert. Springer Lecture Notes in Computer Science 1144. 1996.
4. M.D. Greenberg. *Foundations of Applied Mathematics*. Prentice-Hall Inc. Englewood Cliffs, NJ. 1978.
5. P.W. Hallinan. "A low-dimensional lighting representation of human faces for arbitrary lighting conditions". In *Proc. IEEE Conf. on Comp. Vision and Patt. Recog.*, pp 995-999. 1994.
6. P.W. Hallinan. *A Deformable Model for Face Recognition under Arbitrary Lighting Conditions*. PhD Thesis. Division of Applied Sciences. Harvard University. 1995.
7. B.K.P. Horn. **Computer Vision**. MIT Press, Cambridge, Mass. 1986.
8. B.K.P. Horn and M. J. Brooks, Eds. **Shape from Shading**. Cambridge MA, MIT Press, 1989.
9. P.J. Huber. *Robust Statistics*. John Wiley and Sons. New York. 1981.
10. S.K. Nayar, K. Ikeuchi and T. Kanade "Surface reflections: physical and geometric perspectives" *IEEE trans. on Pattern Analysis and Machine Intelligence*, vol 13 p611-634. 1991.
11. J. Oliensis. "Shape from shading as a partially well-constrained problem". *Computer Vision, Graphics, and Image Processing: Image Understanding*. **54**, pp 163-183. 1991.
12. C. Von Westenholz. **Differential forms in Mathematical Physics**. North-Holland publishing company. Amsterdam. 1981.
13. A.L. Yuille, M. Ferraro, and T. Zhang. "Surface Shape from Warping and Object Recognition". Smith-Kettlewell Eye Research Institute Preprint. 1996.