LOCAL SYMMETRIES OF SHAPES IN ARBITRARY DIMENSION

1. Introduction

In [TSP], level curves of a function v, called "the edge strength function," defined for 2–dimensional shapes, are interpreted as successively smoother versions of the initial shape boundary. The local minima of the absolute gradient $||\nabla v||$ along the level curves of v are shown to be a robust criterion for determining the shape skeleton. More generally, at an extremal point of $||\nabla v||$ along a level curve, the level curve is locally symmetric with respect to the gradient vector ∇v. That is, at such a point, the level curve is approximately a conic section whose one of the principal axes coincides with the gradient vector. Thus, the locus of the extremal points of $||\nabla v||$ along the level curves determines the axes of local symmetries of the shape. In this paper, we extend this method to shapes of arbitrary dimension and illustrate it by applying to 3D shapes.

2. The Edge-Strength Function

Let Ω be a connected, bounded, open subset in $\Re^n$ representing an n-dimensional shape. Let Γ be the boundary of Ω. We consider the following functional:

(1) $$\Lambda_\rho(v) = \frac{1}{2}\int\limits_\Omega \left\{ \rho\|\nabla v\|^2 + \frac{v^2}{\rho} \right\}$$

subject to the boundary condition $v = 1$ along Γ. Let v denote the unique minimizer of the functional. Then, v varies between 0 and 1 and decays exponentially away from Γ. As $\rho \to 0$, $v \to 0$ everywhere except along Γ. Thus, v may be thought of as a blurred version of the characteristic function of Γ and ρ as the nominal blurring radius. This functional was introduced by Ambrosio and Tortorelli [AT] in the context of approximating the Mumford-Shah segmentation functional and in that context v may be interpreted as the probability for the presence of an edge. For this reason, we call v the edge-strength function. The key point is that as $\rho \to 0$, $\min \Lambda_\rho(v)$ tends to the "volume" of Γ. We compute the minimizer of functional (1) by numerically computing the steady state of the following linear diffusion equation obtained by applying gradient descent to functional (1):

(2) $$\frac{\partial v}{\partial \tau} = \nabla^2 v - \frac{v}{\rho^2}$$

The equation may be implemented by means of central finite differences. Alternatively, one can directly solve the steady state equation $\nabla^2 v = v/\rho^2$, by the finite element method for example.

2. 2–Dimensional Shapes

When $n = 2$, $\frac{d||\nabla v||}{ds} = v_{\eta\xi}$ where η is in the direction of the gradient vector ∇v, ξ is in the direction tangent to the level curve, s is the arc-length along the level curve and the subscripts indicate derivatives with respect to these variables. In terms of the global coordinates x, y,

$$v_{\eta\xi} = \frac{\left\{ \left(v_y^2 - v_x^2 \right) v_{xy} - v_x v_y (v_{yy} - v_{xx}) \right\}}{||\nabla v||^2} \tag{3}$$

Hence, the extremal points of $||\nabla v||$ along the level curves of v are given by the zero-crossings of $v_{\eta\xi}$. The symmetry of the level curve at a point P where $v_{\eta\xi} = 0$ is indicated by the missing $\eta\xi$–term in the Taylor series of v at P:

$$v = a_{00} + a_{10}\eta + a_{20}\eta^2 + a_{02}\xi^2 + \cdots \tag{4}$$

Thus, locally at P, the level curve $v = a_{00}$ is approximately a conic section whose one of the principal axes coincides with the gradient vector. An equivalent description of the symmetry at P is that the hessian of v at P is diagonalized when expressed in local coordinates η and ξ. This means that the gradient vector ∇v is an eigenvector of the hessian at P, the other eigenvector being tangent to the level curve at P.

3. General Case: n arbitrary

We now extend our analysis of the 2–dimensional case to arbitrary dimension. First, we look for the points where $||\nabla v||$ is stationary along the level hypersurfaces of v.

Proposition: $||\nabla v||$ is stationary along a level hypersurface at a point P if and only if ∇v is an eigenvector of the hessian H of v.

Proof: $||\nabla v||$ is stationary along a level hypersurface at a point P if and only if the derivative of $||\nabla v||$ in any direction tangent to the level hypersurface at P vanishes. That means that at P, $\nabla||\nabla v||$ cannot have a component tangent to the level hypersurface at P. In other words, the directions of $\nabla||\nabla v||$ and ∇v must coincide at P. That is, $\nabla||\nabla v||$ must be a multiple of ∇v. But

$$\nabla(||\nabla v||) = \frac{H\nabla v}{||\nabla v||} \tag{5}$$

It follows that the necessary and sufficient condition for P to be a stationary point is that

$$H\nabla v = c\nabla v \text{ for some constant } c. \tag{6}$$

Q.E.D.

If ∇v is the vector $\{v_1, v_2, ..., v_n\}$ and $H\nabla v$ is the vector $\{w_1, w_2, ..., w_n\}$, then we have $n-1$ equations

$$\frac{w_1}{v_1} = \frac{w_2}{v_2} = \cdots = \frac{w_n}{v_n} \tag{7}$$

to determine the *1–dimensional* locus of the extremal points. Note that the hessian at an extremal point is diagonalized if we choose the direction of the gradient vector as one of the local coordinates and choose the other coordinates appropriately in the hyperplane tangent to the level hypersurface. Thus the convexity of v is symmetric with respect to the direction of ∇v.

To obtain more information about the shape, we now look for partial symmetries. That is, we require that the level hypersurface be symmetric only about some linear space containing ∇v. To be more specific, consider the matrix Σ whose columns are

$$\nabla v, H\nabla v, H^2\nabla v, \cdots, H^{n-1}\nabla v \tag{8}$$

Observe that if ∇v is an eigenvector of H, the rank of Σ is one. Therefore we define the partial symmetry locus S_k of dimension k as the locus of points where Σ has rank $\leq k$. We now have a sequence of nested loci of successively increasing degree of symmetry:

$$\Omega = S_n \supset S_{n-1} \supset \cdots \supset S_1 \supset S_0 \tag{9}$$

where S_0 is just the locus of points where ∇v vanishes, making the rank of Σ equal to zero. At a point in S_k, v is locally symmetric with respect to the linear space spanned by the columns of Σ. It contains the gradient vector ∇v and is spanned by k eigenvectors of H. Note that our definition of S_{n-2} is analogous to the definition of Furst, Pizer and Eberly [MMBIA] for the ridges of a function defined over a 4–dimensional domain.

4. 3–Dimensional Shapes

In the case of 3–dimensional shapes, S_2 is given simply by the vanishing of the determinant of the matrix Σ. The 1–dimensional locus S_1 is given by the equations

(10)
$$\frac{v_x v_{xx} + v_y v_{xy} + v_z v_{xz}}{v_x} = \frac{v_x v_{xy} + v_y v_{yy} + v_z v_{yz}}{v_y} = \frac{v_x v_{xz} + v_y v_{yz} + v_z v_{zz}}{v_z}$$

which may be written as

(11)
$$\begin{aligned} v_x v_y(v_{xx} - v_{yy}) + v_{xy}\left(v_y^2 - v_x^2\right) + v_z(v_y v_{xz} - v_x v_{yz}) &= 0 \\ v_y v_z(v_{yy} - v_{zz}) + v_{yz}\left(v_z^2 - v_y^2\right) + v_x(v_z v_{yx} - v_y v_{zx}) &= 0 \\ v_z v_x(v_{zz} - v_{xx}) + v_{zx}\left(v_x^2 - v_z^2\right) + v_y(v_x v_{zy} - v_z v_{xy}) &= 0 \end{aligned}$$

The 0–dimensional locus S_0 is given by the equations $v_x = v_y = v_z = 0$.

5. Shape-Skeletons

Consider first the 2–dimensional case. There are two definitions given in [TSP] for extracting the skeleton of a two-dimensional shape from S_1. The one simplest to implement defines the skeleton as consisting of S_0 and those points of $S_1 \backslash S_0$ where the level curve has positive curvature. That is, the relevant points are where the second derviative of v in the direction orthogonal to ∇v is positive. The points where curvature is negative indicate the presence of a neck. Hence for the general case, we look for a measure which depends only on the second derivatives of v along directions orthogonal to the linear space spanned by the columns of Σ. In the case of S_1, such a measure is readily provided by the mean curvature μ given by the formula

(12)
$$\begin{aligned} \mu &= \nabla \cdot \frac{\nabla v}{||\nabla v||} \\ &= \frac{1}{||\nabla v||}\left[\nabla \cdot \nabla v - \nabla^2 v\left(\frac{\nabla v}{||\nabla v||}, \frac{\nabla v}{||\nabla v||}\right)\right] \end{aligned}$$

where $\nabla^2 v$ is the hessian H of v and the last term is just the second derivative of v in the direction of ∇v. The term $\nabla \cdot \nabla v$ is of course the laplacian of v given by the trace of H. We prune $S_1 \backslash S_0$ by removing points where μ is negative. To extend this construction to every S_k, note that the expression in the bracket is just the sum of the eigenvalues corresponding to the eigenvectors of H orthogonal to ∇v. Therefore, we consider the linear space L spanned by the columns of Σ and the linear space T orthogonal to L. Choose a basis for $\mathfrak{R}^{\mathfrak{n}}$ by choosing an orthonormal basis for L and an orthonormal basis for T. Then at points in $S_k \backslash S_{k-1}$, H is block-diagonal with respect to this basis, consisting of a square block Λ_k of dimension k and a square block Θ_k of dimension $n-k$. At points of $S_k \backslash S_{k-1}$, define

$$\mu_k = \frac{trace(\Theta_k)}{||\nabla v||} \tag{13}$$

The numerator in the expression for μ_k is simply the sum of the second derivatives of v along the basis vector of T. Prune S_k by removing points where μ_k is negative. Note that μ_1 is the mean curvature μ.

In the 3D case, the explicit formula for the mean curvature μ_1 is as follows:

$$\begin{aligned} ||\nabla v||^3 \mu = {} & \left(v_y^2 + v_z^2\right) v_{xx} + \left(v_z^2 + v_x^2\right) v_{yy} + \left(v_x^2 + v_y^2\right) v_{zz} \\ & - 2 v_x v_y v_{xy} - 2 v_y v_z v_{yz} - 2 v_x v_z v_{xz} \end{aligned} \tag{14}$$

To calculate μ_2, let t be the unit vector given by

$$t = \frac{\nabla v \times H \nabla v}{||\nabla v \times H \nabla v||} \tag{15}$$

where $\times$ indicates the cross-product. Then in the 3D case,

$$\mu_2 = \frac{v_{tt}}{||\nabla v||} = \frac{t \cdot Ht}{||\nabla v||} \tag{16}$$

6. EXAMPLES