Lecture 11: Differential Geometry

Last modified on February 28, 2000 at 8:45 PM

Contents

Reading

These notes.

11.1 Introduction

11.1.1 What Differential Geometry Is

Differential Geometry is very much what the name implies: geometry done using differential calculus. In other words, shape description through derivatives.

11.1.2 Assumptions of Differentiability

The first major assumption made in differential geometry is, of course, that the curve, surface, etc. is everywhere differentiable—there are no sharp corners in a differential-geometry world. (Of course, once one accepts the notion of scaled measurements, there are no infinitely-precise corners anyway.)

11.1.3 Normals and Tangent Planes

The first obvious derivative quantity is the curve or surface *tangent*. For curves, this is a vector; for surfaces, this is a *tangent plane*. For differential geometry of 3D shapes, there is no "global" coordinate system. All measurements are made relative to the local tangent plane or normal.

For images, though, we're going to use a coordinate system that defines the image intensity as the "up" direction.

Be aware that differential geometry as a means for analyzing a function (i.e., an image) is quite different from differential geometry on general surfaces in 3D. The concepts are similar, but the means of calculation are different.

11.1.4 Disclaimer

If a true differential geometer were to read these notes, he would probably cringe. This is a *highly* condensed and simplified version of differential geometry. As such, it contains no discussion of forms (other than the Second Fundamental Form), covectors, contraction, etc. It also does not attempt to address non-Euclidean aspects of differential geometry such as the bracketing, the Levi-Civita tensor, etc.

11.1.5 Notation

For this lecture, we'll use the notation common in most of the literature on the subject. This notation uses subscripts to denote derivatives in the following fashion.

If we let $L(x, y)$ be our image's luminance function (pixel values), the derivative with respect to x is L_x, and the derivative with respect to y is L_y. In fact, we can denote the derivative of L with respect to *any* direction (unit vector) $\bar{v}$ as L_v.

By extension, the second derivative with respect to x is L_{xx}, the second derivative with respect to y is L_{yy}, and the mixed partial (the derivative of L_x with respect to h is L_{xy}. In general, the derivative with respect to vector $\bar{u}$ *as it changes in direction* $\bar{v}$ is L_{uv}.

11.2 First Derivatives

11.2.1 The Gradient

The primary first-order differential quantity for an image as the *gradient*. The gradient of a function f is defined as

$$\nabla L(x, y) = \begin{bmatrix} L_x(x, y) \\ L_y(x, y) \end{bmatrix}$$

Notice that the gradient is a 2-D vector quantity. It has both direction and magnitude *which vary at each point*. For simplicity, we usually drop the pixel locations and simply write

$$\nabla L = \begin{bmatrix} L_x \\ L_y \end{bmatrix}$$

The gradient at a point has the following properties:

- The gradient direction at a point is the direction of steepest ascent at that point.
- The gradient magnitude is the steepness of that ascent.
- The gradient direction is the normal to the level curve at that point.

- The gradient defines the tangent plane at that point.
- The gradient can be used to calculate the first derivative *in any direction*:

$$L_v = \bar{v} \cdot \nabla L$$

The gradient can thus be used as a *universal first-derivative calculator*. It includes all first-derivative information there is at a point.

11.2.2 First-order Gauge Coordinates

One of the fundamental concepts of differential geometry is that we can describe local surface properties not with respect to some global coordinate system but with respect to a coordinate system *dictated by the local (image) surface itself*. This concept is known as *gauge coordinates*, a coordinate system that the surface "carries along" with itself wherever it goes. Many useful image properties can likewise be most easily described using gauge coordinates. (More on this later.)

The gradient direction is one of those intrinsic image directions, independent of our choice of spatial coordinate axis. (The numbers we use to express the gradient as a vector in terms of our selected coordinate system will, of course, change as we change coordinates, but the gradient direction itself will not.)

The gradient direction and its perpendicular constitute *first-order gauge coordinates*, which are best understood in terms of the image's level curves. An *isophote* is a curve of constant intensity. The normal to the isophote curve is the gradient direction. Naturally, the tangent to the isophote curve is the gradient perpendicular.

We denote the first-order gauge coordinates as directions v and w where

$$\begin{aligned} w &= \frac{\nabla L}{\|\nabla L\|} \\ v &= w_{\perp} \end{aligned}$$

Local properties of isophotes (image level curves) are best described in terms of these gauge coordinates.

11.3 Second Derivatives

11.3.1 The Hessian

For second-order geometry, the equivalent of the gradient is the matrix of second derivatives or *Hessian*:

$$H = \begin{bmatrix} L_{xx} & L_{xy} \\ L_{yx} & L_{yy} \end{bmatrix}$$

Since $L_{xy} = L_{yx}$, this matrix is symmetric.

We can use the Hessian to calculate second derivatives in this way:

$$L_{uv} = \bar{u}^T H \bar{v}$$

or if we use the same vector on both sides of the matrix:

$$L_{vv} = \bar{v}^T H \bar{v}$$

In other words, it's a sort of "universal second derivative calculator"

Here's an example. The unit vector in the x direction is $[1, 0]^T$. We get the second derivative in this direction by calculating

$$\begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} L_{xx} & L_{xy} \\ L_{yx} & L_{yy} \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

This is simply L_{xx}. We can do this with any unit vector in any direction.

The Hessian matrix, being real and symmetric, has a few interesting properties:

1. Its determinant is equal to the product of its eigenvalues and is invariant to our selection of x and y.
2. Its trace (sum of the diagonal elements) is also invariant to our selection of x and y.
3. Proper selection of x and y can produce a diagonal matrix—these are the eigenvectors of H, and the diagonal elements are the eigenvalues.

11.3.2 Principal Curvatures and Directions

The eigenvalues and eigenvectors of the Hessian have geometric meaning:

- The first eigenvector (the one whose corresponding eigenvalue has the largest absolute value) is the direction of greatest curvature (second derivative).
- The second eigenvector (the one whose corresponding eigenvalue has the smallest absolute value) is the direction of least curvature.
- The corresponding eigenvalues are the respective amounts of these curvatures.

The eigenvalues of H are called *principal directions* and are directions of pure curvature (no mixed partial derivative). They are always orthogonal.

The eigenvalues of H are called *principal curvatures* and are invariant under rotation. They are denoted κ_1 and κ_2 and are always real valued.

11.3.3 Second-order Gauge Coordinates

Just as we used the gradient direction and perpendicular as *first-order gauge coordinates*, we can use the principal directions as *second-order gauge coordinates*. These second-order gauge coordinates, like their first-order counterparts, are intrinsic to the image surface. They are the most useful (and most natural) coordinate system for describing second-order properties.

We denote these second-order gauge directions (the principal curvature directions) as p and q. So, $L_{pp} = \kappa_1$, $L_{qq} = \kappa_2$, and by convention $|L_{pp}| \geq |L_{qq}|$.

11.4 Using Principal Curvatures and Directions

11.4.1 Gaussian Curvature

We can write the determinant of H as the product of κ_1 and κ_2. This quantity is the *Gaussian Curvature* and is denoted as K.

Geometrically, it too has a meaning. Suppose that you took a local patch around a point on the surface and “squashed” it flat. To do so, you’d have to tear the surface, because the patch has more area than its perimeter would normally contain for a flat region. The Gaussian curvature is the amount of this “extra stuff”.

In second-order gauge coordinates, the Gaussian curvature is $L_{pp}L_{qq}$.

11.4.2 Mean Curvature

The mean curvature is the average of κ_1 and κ_2 and is denoted as H. It is also equal to the half the trace of H, which we earlier said was invariant to our selection of x and y.

In second-order gauge coordinates, the mean curvature is $(L_{pp} + L_{qq})/2$.

11.4.3 Laplacian

The *Laplacian* that you learned about in CS 450 can be written as $L_{xx} + L_{yy}$. This is simply twice the mean curvature and is also invariant to rotation.

11.4.4 Deviation From Flatness

The *deviation from flatness* is $L_{pp}^2 + L_{qq}^2$ and is another useful way of measuring local "unflatness".

11.4.5 Patch Classification

Surface regions can be classified according to their mean and Gaussian curvatures.

- Elliptic patches: $K > 0$. The curvature in any direction is positive.
  - Convex: $H > 0$
  - Concave: $H < 0$. The curvature in any direction is negative.
- Hyperbolic patches: $K < 0$. The curvatures in some directions are positive and others are negative.

11.4.6 Parabolic Points

We've covered the cases for $K > 0$ and $K < 0$. If K exactly equals 0, one or both of the principal curvatures is 0. Such points form curves on the surface known as *parabolic curves*. These curves naturally lie on the boundaries between elliptic and hyperbolic regions.

11.4.7 Umbilics

An *umbilic* is a point whose principal curvatures are equal: $\kappa_1 = \kappa_2$. At such a point, the eigenvectors are not uniquely defined (the eigenspace of H is the entire tangent plane). In other words, such a point has the same curvature *in every direction* (locally spherical). Since the principal directions are not defninted at an umbilic, we generally choose them to vary smoothly through the umbilic.

11.4.8 Minimal Points

Hyperbolic regions are saddle-shaped. One special point that can occur in such regions is the cousin of an umbilic: a point with equal-magnitude but opposite-sign principal curvatures: $\kappa_1 = -\kappa_2$. Such a point is known as a *minimal point*. (The name comes from *minimal surfaces*, the soap-bubble surfaces that form across holes by minimizing the surface area across the film.)

11.4.9 Shape Classification

Jan Koenderink has proposed a "shape classification" space where each point lives in a Cartesian $\kappa_1 \times \kappa_2$ plane. However, this space can best be thought of in its polar form:

$$
\begin{aligned}
S &= \tan^{-1}\left(\frac{\kappa_2}{\kappa_1}\right) \\
C &= \sqrt{\kappa_1^2 + \kappa_2^2}
\end{aligned}
$$

where S defines the shape angle, and C defines the degree of curvature. (Observe that C is the square-root of the deviation from flatness.) Points with the same S value but differing C values can be thought of as being the same shape but only "more stretched".

Quadrant I of this space is convex, Quadrant III is concave, and Quadrants II and IV are hyperbolic. The Cartesian axes of the space are parabolic points. At 45 and -135 degrees lie the umbilics; at 135 and -45 degrees lie the minimal points.

11.5 Other Concepts

11.5.1 Geodesics

A *geodesic* is a path of shortest distance between two points. On a plane, geodesics are straight lines ("the shortest path between two points is a straight line"). On a surface, though, this isn't necessarily true. In fact, there may be multiple geodesics between two points. "Great circles" on (flattened) earth maps are an example of geodesics.

11.5.2 Genericity

Points with special properties are usually only important if they are predictable and stable instead of coincidental. Properties that occur by coincidence go away with a small perturbation of the geometry. Consider, for example, two lines intersecting in three-space—their intersection goes away if either line is perturbed by a small amount. Such situations are called *nongeneric*.

Consider, however, the intersection of two lines in a two-dimensional plane. If either line moves a small amount, the point of intersection may move, but it doesn't go away. Such situations are called *generic*.

11.6 Applications (Examples)

11.6.1 Corners and Isophote Curvature

Corners are places where the isophote has high curvature. The isophote curvature can be written as

$$\kappa = -\frac{L_{vv}}{L_w}$$

Homework exercise: Using what we've talked about so far, express this in terms of L_x, L_y, L_{xx}, L_{yy}, and L_{xy}.

11.6.2 Ridges

Ridges are an intuitively simple concept, but their mathematical definition is not well agreed upon. One definition is that a ridge is a maximum of image intensity in the direction of maximal curvature. This may be written as two constraints:

1. $L_p = 0$
2. $L_{pp} < 0$

This is simply the classic maximum test from calculus.

Another is that a ridge is a maximum of isophote curvature (i.e., a ridge is a connected locus of isophote corners). This can be defined as

$$L_{vvv} = 0$$

11.7 Conclusion

The whole aim of this discussion of differential geometry on image surfaces is to introduce the notion of gauge coordinates as a way of measuring *invariants*, quantities that do not change with arbitrary selection of spatial coordinate systems. The following quantities are invariant to rotation:

- gradient magnitude L_w (L_v is also invariant, but it's always zero)
- the first principal curvature L_{pp}
- the second principal curvature L_{qq}

The following quantities can be computed from these and are thus also invariant:

- isophote curvature
- the Laplacian
- Gaussian curvature
- mean curvature
- deviation from flatness
- all special points such as umbilics, minimal points, ridges, corners, etc.

Vocabulary

- Gradient
- Hessian
- Gauge coordinates
- First-order gauge coordinates
- Second-order gauge coordinates
- Principal directions
- Principal curvatures
- Gaussian curvature
- Mean curvature
- Deviation from flatness
- Laplacian
- Elliptical regions
- Hyperbolic regions
- Parabolic point/curve
- Umbilic
- Minimal points
- Shape classification
- Isophote
- Isophote curvature
- Ridge
- Generic geometric properties