Lectures in Discrete Differential Geometry 2 – Surfaces

Etienne Vouga

February 14, 2014

1 Smooth Surfaces in $\mathbb{R}^3$

In this section we will review some properties of smooth surfaces $M \subset \mathbb{R}^3$. We will assume that M is parameterized by a single chart $r(u,v) : \tilde{M} \to \mathbb{R}^3$, where $\tilde{M} \subset \mathbb{R}^2$. (Surfaces of practical interest, including basic examples like the round sphere or torus of revolution, cannot be parameterized in this way, and instead are usually described using an atlas of multiple such charts "glued together" in a compatible way at their regions of overlap. See any introductory text on differential geometry for more details; we will ignore this issue in what follows.)

As for curves, we will assume that M is well-behaved: that r is C^∞ and an immersion, i.e. that the partial derivatives r_u and r_v do not vanish and are linearly independent at every point $\tilde{x} = (u,v) \in \tilde{M}$. (As a helpful notational convention, throughout this section we will label quantities on the plane with tildes, and quantities in space without.)

As was the case with plane curves, we are interested in the geometric quantities we can measure on M. Some first examples include

- *Tangent planes* at points on M. At every point $x = r(\tilde{x})$ of M, the tangent plane $T_x M$ is spanned by $r_u(\tilde{x})$ and $r_v(\tilde{x})$. Tangent vectors w at x on M can be mapped to tangent vectors $\tilde{w}$ at $\tilde{x}$ in the plane, and vice versa, in a canonical way, using the directional derivative of curves. Consider the curve $\tilde{\gamma}(t) = \tilde{x} + t\tilde{w}$ in $\mathbb{R}$. The parameterization r maps this curve to a curve $\gamma(t) = r \circ \tilde{\gamma}(t)$ on M. The tangent vector $\tilde{w}$ of $\tilde{\gamma}$ at $t = 0$ then maps to the tangent vector $\gamma'(0)$, which by the chain rule is equal to $dr_{\tilde{x}}(\tilde{w}) = [dr_{\tilde{x}}]\,\tilde{w}$, where $dr_{\tilde{x}}$ is the 3×2 Jacobian $[r_u(\tilde{x})\ r_v(\tilde{x})]$ of r.

- *Lengths and angles* of tangent vectors, as measured using the ordinary Euclidean dot product $w_1 \cdot w_2$ for any pair of tangent vectors w_1, w_2 at $x \in M$. Since these tangent vectors can be represented using vectors $\tilde{w_1}, \tilde{w_2}$ at $\tilde{x} \in \mathbb{R}^2$, it follows that it must be possible to measure $w_1 \cdot w_2$ using $\tilde{w_1}$ and $\tilde{w_2}$... but not in the most naive way! Generally $w_1 \cdot w_2 \neq \tilde{w_1} \cdot \tilde{w_2}$. Instead, length and angles on M can be measured on $\tilde{M}$ by using the fact that tangent vectors on M and $\tilde{M}$ are related by the Jacobian of r:

$$w_1 \cdot w_2 = dr(\tilde{w_1}) \cdot dr(\tilde{w_2}) = dr(\tilde{w_1})^T dr(\tilde{w_2}) = \tilde{w_1}^T \left[r_u\ r_v\right]^T \left[r_u\ r_v\right] \tilde{w_2} = \tilde{w_1}^T g \tilde{w_2},$$

  where

$$g(\tilde{x}) = \begin{bmatrix} r_u(\tilde{x}) \cdot r_u(\tilde{x}) & r_u(\tilde{x}) \cdot r_v(\tilde{x}) \\ r_u(\tilde{x}) \cdot r_v(\tilde{x}) & r_v(\tilde{x}) \cdot r_v(\tilde{x}) \end{bmatrix}$$

  is the *metric* of the surface (or, more formally, the metric on $\tilde{M}$ that you get by pulling back the Euclidean metric on space using r.) The metric g is sometimes also called the *first fundamental form* of M, as it encodes all intrinsic information about M – that is, if you know $\tilde{M}$ and g, but nothing else about r, you know everything that depends only on lengths, distances, and angles on M. For instance, you can calculate the arc length of any curve $\gamma(t) = r \circ \tilde{\gamma}(t) : [0,1] \to M$:

$$L(\gamma) = \int_0^1 \|\gamma'(t)\|\,dt = \int_0^1 \|dr(\tilde{\gamma}')\|\,dt = \int_0^1 \sqrt{\tilde{\gamma}'^T g \tilde{\gamma}'}\,dt.$$

- *Normal vectors* $N(x)$ at points $x \in M$:

$$N(x) = \frac{r_u(\tilde{x}) \times r_v(\tilde{x})}{\|r_u(\tilde{x}) \times r_v(\tilde{x})\|}.$$

  Notice that N maps neighborhoods of M to neighborhoods of the unit sphere – this map $N(x)$ is often called the *Gauss map*.

- *Areas.* Infinitesimal area $dudv$ gets mapped to an infinitesimal parallelogram with sides dur_u and dvr_v; therefore

$$dA = \|r_u \times r_v\| \, d\tilde{A} = \sqrt{\|r_u\|^2 \|r_v\|^2 - (r_u \cdot r_v)^2} \, d\tilde{A} = \sqrt{\det g} \, d\tilde{A},$$

  which allows us to integrate functions over the surface M by integrating over the appropriate region of the plane instead.

Functions $f(x) = f(r(\tilde{x})) : M \to \mathbb{R}$ have a natural notion of directional derivative on M in the tangent direction w:

$$D_w f = \frac{d}{dt} f(r(\tilde{x} + t\tilde{w}))\bigg|_{t \to 0} = \frac{d}{dt} \tilde{f}(\tilde{x} + t\tilde{w})\bigg|_{t \to 0} = \nabla \tilde{f} \cdot \tilde{w},$$

where $\nabla \tilde{f}$ is the ordinary gradient in $\mathbb{R}^2$. The inner product

$$\langle f, g \rangle = \int_M f g \, dA$$

on functions on M induces an *intrinsic* gradient ∇f: the tangent vector field on M with the property that for every tangent vector field δ on M,

$$\langle \nabla f, \delta \rangle = \int_M \nabla f \cdot \delta \, dA = D_\delta f.$$

1.1 Curvatures

We know how to measure the curvature of curves, and can use this to understand curvature of surfaces. Pick a point $x \in M$ and tangent direction w at x, and *cut* the surface with the plane spanned by w and $N(x)$. The curvature of the resulting curve, at x, is the *normal curvature* $\kappa_w(x)$.

Let $\gamma(t)$ be the curve formed by the cut. From the last lecture, we know that

$$\kappa(s) = -T'(s) \cdot N(s) = -T'(s)^T J T(s) = T'(s)^T J^T T(s) = N'(s)^T T(s),$$

and if the curve is not parameterized by arc length, the equivalent formula is, by the chain rule,

$$\kappa(t) = \frac{N'(t) \cdot T(t)}{\|\gamma'(t)\|} = \frac{N'(t) \cdot \gamma'(t)}{\|\gamma'(t)\|^2}.$$

Normal curvature is then

$$\kappa_w(x) = \frac{D_{\gamma'(0)} N \cdot \gamma'(0)}{\|w\|^2} = \frac{D_w N \cdot w}{\|w\|^2}.$$

Since N is a unit vector, its derivative, in any tangent direction, must be orthogonal to it and so a tangent vector of N. Therefore there exists a map $S : T_x M \to T_x M, w \mapsto D_w N$ from tangent vectors at x to tangent vectors at x, and since the directional derivative is linear in the direction w, S is also a linear operator, which can be represented as a 3×3 matrix. A linear map on w is also a linear map on $\tilde{w}$, so there exists a 2×2 matrix $\tilde{S}$ with $dr(\tilde{S}\tilde{w}) = Sw$, so that the normal curvature, either in space or in coordinates, is given by

$$\kappa_w(x) = \frac{w \cdot Sw}{\|w\|^2} = \frac{\tilde{w}^T g \tilde{S} \tilde{w}}{\tilde{w}^T g \tilde{w}}.$$

The map S is called the *shape operator*, or sometimes the *Weingarten map*, and it can be calculated that

$$\tilde{S} = g^{-1}b, \qquad b(\tilde{x}) = \begin{bmatrix} r_{uu}(\tilde{x}) \cdot N(\tilde{x}) & r_{uv}(\tilde{x}) \cdot N(\tilde{x}) \\ r_{uv}(\tilde{x}) \cdot N(\tilde{x}) & r_{vv}(\tilde{x}) \cdot N(\tilde{x}) \end{bmatrix},$$

where the matrix b is called the *second fundamental form* of the surface M. Whereas the first fundamental form, g, encodes all intrinsic information about M, b encodes the rest: if you know it, you know how the surface curves in every direction, and so can completely reconstruct the surface (up to rigid motions).

A few remarks are in order. First, κ_w is scale invariant – both numerator and denominator of the above formulas scale quadratically with the scale of w – so we can take w to be a unit vector $\|w\| = 1$ and thereby ignore the denominator. Second, from the above formula for b it follows that the second fundamental form, or equivalently the shape operator, is symmetric (a fact not obvious from its definition as the derivative of the normal vector). The matrix in coordinates, $\tilde{S}$, is the product of two symmetric matrices and is therefore not generally symmetric itself – it is, however, self-adjoint under the g metric: $\tilde{w_1}^T g\tilde{S}\tilde{w_2} = \tilde{w_1}^T \tilde{S}^T g\tilde{w_2}$ for any pair of vectors $\tilde{w_1}$ and $\tilde{w_2}$. Therefore by the spectral theorem, $\tilde{S}$ has a pair of eigenvectors $\tilde{d_1}, \tilde{d_2}$ with corresponding real eigenvalues λ_1, λ_2, satisfying:

- $\tilde{d_i}^T g\tilde{d_i} = 1$, or equivalently, the three-dimensional tangent vectors $d_i = dr(\tilde{d_i})$ are unit-length;
- $\tilde{S}\tilde{d_i} = \lambda_i \tilde{d_i}$;
- $\tilde{d_1}^T g\tilde{d_2} = 0$, or equivalently, d_1 and d_2 are orthogonal;
- $\lambda_1 \leq \lambda_2$ (without loss of generality); and
- $\lambda_1 \leq \kappa_w \leq \lambda_2$ for any direction w. This follows from decomposing $\tilde{w}$ in the eigenbasis.

The eigenvalues λ_i are called the *principal curvatures* at x and the eigenvectors the *principal curvature directions*. The above properties of principal curvature have strong geometric meaning: the two principal curvature directions are the directions of most and least (signed) curvature, and these directions are always orthogonal. Moreover, the normal curvature in any other direction is completely determined by the principal curvatures, and is somewhere between these two extreme curvatures.

The principal curvatures and directions give the normal curvature of the surface in every direction, but there are also two important measures of curvature at a point that do not depend on choosing a direction. The first of these is the "average" normal curvature

$$\frac{1}{2\pi} \int_{\|w\|=1} k_w(x)$$

that you get by computing $k_w(x)$ along every possible tangent direction w. It can be shown that the above average is equal to $H = \frac{1}{2} \operatorname{tr} \tilde{S} = \frac{\lambda_1 + \lambda_2}{2}$, the *mean curvature* at x. (Danger: some authors define the mean curvature as $\operatorname{tr} \tilde{S} = 2H$. This simplifies several formulas, but does not sensibly reflect the name, and easily causes confusion.)

The other basis-invariant measure of $\tilde{S}$ is the determinant; $K = \det \tilde{S} = \lambda_1 \lambda_2$ is called the Gaussian curvature at x. Its sign qualitatively describes the shape of the surface near x: bowl-like (two equal directions of curvature) when $K > 0$, saddle-like (two opposite directions of curvature) when $K < 0$, and developable (one flat direction) when $K = 0$.

As an example, here are the curvatures of several basic surfaces:

- A plane has zero mean and Gaussian curvature.
- A sphere of radius r has $\lambda_1 = \lambda_2 = \frac{1}{r}$, so $H = \frac{1}{r}, K = \frac{1}{r^2}$.
- A cylinder of radius r has minimum curvature $\lambda_1 = 0$, along the axis of the cylinder, and maximum curvature $\lambda_2 = \frac{1}{r}$ in the circumferential direction. Therefore $H = \frac{1}{2r}$ and $K = 0$.

- A torus has different curvature at different points. Suppose the outer radius of the torus is r_1 and the inner radius r_2. At a point along the "outer equator," $\lambda_1 = \frac{1}{r_1+r_2}$, $\lambda_2 = \frac{1}{r_2}$, and Gaussian curvature is positive. At a point on the "top," $\lambda_1 = 0, \lambda_2 = \frac{1}{r_2}$, and Gaussian curvature is zero. On the "inner equator," $\lambda_1 = \frac{-1}{r_1-r_2}$, $\lambda_2 = \frac{1}{r_2}$, and Gaussian curvature is negative.

1.2 Properties of Curvature

Mean and Gaussian curvature satisfy many of the properties that help for curvature of plane curves. Definition of curvature is different, of course: $H = \frac{1}{2}\operatorname{tr}(g^{-1}b), K = \det(g^{-1}b)$. As for the other properties:

1. H and K are both invariant under transformation of M by rigid motions in space. More than that, though, it turns out that K is invariant under arbitrary isometric deformations of the surface! In other words, even though the definition of K depends on b, it could be rewritten to only depend on the metric g. This is Gauss's celebrated "egregious" (or "noteworthy") theorem.

2. Reflections in $\mathbb{R}^3$ flip the sign of H, but leave K unchanged. This property, as well as the two that follow, are obvious consequences of the relationship between H and K and the principal curvatures.

3. Reversing the orientation of M flips the sign of H but leaves K unchanged.

4. H scales as $\frac{1}{\text{length}}$; K scales as $\frac{1}{\text{length}^2}$.

5. For a family of smooth surfaces converging to a surface with a "kink" (i.e., if $\kappa_w(x)$ diverges at some point x and direction w), then one of mean or Gaussian curvature also diverges (not necessarily both).

6. $H = K = 0$ for planes.

7. For spheres of radius r (oriented outwardly), $H = \frac{1}{r}$ and $K = \frac{1}{r^2}$.

8. Mean and Gaussian curvature are local: their values at a point $r(\tilde{x})$ depend only on r in a neighborhood of $\tilde{x}$.

9. I'm not aware of any particularly enlightening generalization of the osculating circle property to curvatures of surfaces. Osculating quadratic surface, perhaps?

10. For a closed surface M, $\int_M K\, dA = 4\pi n$ for some $n \in \mathbb{Z}$. Furthermore, n has a geometric intepretation: $n = 1 - g$, where g is the genus of the surface (number of holes). This is the *Gauss-Bonnet* theorem.

    The Gauss-Bonnet theorem follows from a more local result. First, we can define a distance function on M as the length of the shortest curve between two points:

    $$d(x_1, x_2) = \inf_{\substack{\gamma:[0,1]\to M \\ 0\mapsto x_1, 1\mapsto x_2}} L(\gamma),$$

    where L is the arc length functional. We can then define a *geodesic disk* $B_r(x)$ around any $x \in M$ as the set of points a distance less than r away from x on M:

    $$B_r(x) = \{y \in M | d(x, y) < r\}.$$

    (For instance, for the special case that M is a plane, $B_r(x)$ is just an ordinary open disk of radius r.) $B_r(x)$ has some preimage $\tilde{B}_r(x)$ in the plane. It is also possible to look at the set of normal vectors at all points inside $B_r(x)$; call this set $B_r^N(x) = N(\tilde{B}_r(x))$. Then

    $$K(x) = \lim_{r\to 0} \frac{A\left(B_r^N(x)\right)}{A\left(B_r(x)\right)},$$

where the surface area functional A measures the oriented surface area of B_r^N and B_r on the unit sphere or on M, respectively. For example, for a plane, $B_r^N(x)$ contains just a single point on the unit sphere, which has zero area. For a cylinder, B_r^N is a line which still has zero surface area. For a sharply-curved bowl-shaped surface, B_r^N covers a large part of the unit sphere even for small r; for a saddle-shaped surface, B_r^N also covers a patch of the unit sphere, but does so with negative orientation.

Another way of expressing the above identity is as

$$K(x) = \frac{dA^N(x)}{dA(x)},$$

where dA is infinitesimal surface area on M, and dA^N is infitesimal surface area of the unit sphere covered by the normals in an infinitesimal neighborhood of x. Then

$$\int_M K\,da = \int_M dA^N$$

measures, intuitively, the number of times the normals of M "paint" the unit sphere. Since the normal vector changes continuously over M, this must be an integer: the normals of a sphere-like object paint the unit sphere exactly once, for instance, whereas the normals of a torus paint the sphere a total of zero times (positive one times around the outside of the torus, and negative one times over the inside of the torus.)

11. Mean curvature times the normal vector is the gradient of surface area:

$$(\nabla A)(x) = 2H(x)N(x).$$

As for curves in the plane, this is the gradient in the sense of the calculus of variations: for closed surfaces, the vector field $2HN$ over M is the unique vector field with the property that for every vector field δ over M,

$$\langle 2HN, \delta\rangle = \int_M 2HN \cdot \delta\, dA = D_\delta A,$$

where A is the surface area functional and the scalar $D_\delta A$ is the directional derivative

$$D_\delta A = \frac{d}{dt} A(r + t\delta)\bigg|_{t\to 0}.$$

12. Lastly, we have an analogue to the "inflation theorem" for space curves. Let M be a closed, injective surface (one without self-intersections), and $V(M)$ the volume functional that measures the volume enclosed by M. For sufficiently small ϵ, we can "inflate" M by ϵ along its normals to get a new surface $M_\epsilon = \mathrm{im}(r + \epsilon N)$, and look at the volume $V(M_\epsilon)$ enclosed by the inflated surface. Obviously, to zeroth order, this volume is just the original volume $V(M)$, and the first-order term is the surface area of M. As might be expected from the earlier theorem we examined on inflating plane curves, the higher-order terms depend on the integrated curvature of M; dimensional analysis might suggest the following remarkably beautiful formula, the *Steiner polynomial*:

$$V(M_\epsilon) = V(M) + \epsilon A(M) + \epsilon^2 \int_M H\,dA + \frac{\epsilon^3}{3}\int_M K\,dA.$$

For example, the sphere of radius r has inflated volume

$$\frac{4}{3}\pi(r+\epsilon)^3 = \frac{4}{3}\pi r^3 + \epsilon 4\pi r^2 + \epsilon^2(4\pi r^2)\frac{1}{r} + \frac{\epsilon^3}{3}(4\pi r^2)\frac{1}{r^2},$$

as predicted by the Steiner polynomial.

As they did in the case for curves, the above properties will guide our formulation of a discrete notion of curvature for discrete surfaces. In this whirlwind review, we have looked at those properties and formulas that will be most useful in that endeavor; a more complete exposition of surface geometry, including formal proofs of the above properties, can be found in textbooks on Riemannian geometry.