

Lectures in Discrete Differential Geometry 3 – Discrete Surfaces

Etienne Vouga

March 19, 2014

1 Triangle Meshes

We will now study *discrete* surfaces and build up a parallel theory of curvature that mimics the structure of the smooth theory. First, we need a definition of a discrete surface. There are many possible discrete representations in common use – in this course we will focus on *triangle meshes*, though much could be said about alternatives, such as point clouds, quad meshes, tensor product splines, subdivision surfaces, etc.

We define a (manifold, closed) triangle mesh M as a simplicial 2-complex locally homeomorphic to the Euclidean disk. M being a simplicial 2-complex means it consists of

- distinct *vertices* $V = \{v_i \in \mathbb{R}^3\}$;
- oriented *edges* $E = \{e_{ij} \in V \times V\}$;
- oriented *faces* $F = \{f_{ijk} \in V^3\}$

where all boundaries of every face is in E , i.e., if f_{ijk} is a face, then either e_{ij} or e_{ji} is in E , and similarly for the other two sides of the face. Notice that each edge has two possible orientations, as does each face – face orientations are sometimes labeled "clockwise" or "counterclockwise" but of course it depends from which angle you're viewing the face! Notice also that the orientation of a face has nothing to do with the orientations of the edges bounding it.

M being locally homeomorphic to the Euclidean disk means that at each vertex of M , the faces neighboring M "look," topologically, like a piece of the plane. This means that each edge is shared by exactly two faces, faces around a vertex connect in a "fan," and faces have compatible orientation (both clockwise, or both counterclockwise) across edges (see figure 1, left). This means no Klein bottles, Möbius strips, etc. Unless otherwise specified, we will assume that all meshes are closed (no boundaries) and have nondegenerate faces (all faces have noncolinear vertices). The former greatly simplifies the derivations and exposition in the remainder of this lecture; in practice we *do* deal with meshes with boundary all of the time, but since all of the discrete geometric quantities we will derive here will be *local*, there is no difficulty generalizing them




Figure 1: *Left*: Each vertex neighborhood of a triangle mesh looks topologically like an oriented piece of the plane. *Right*: A face and its vertices and edges.

to open meshes (with the possible exception of what to do at mesh boundaries). The latter is the discrete generalization of the requirement that a smooth surface be immersed.

In what follows we will abuse tilde (\sim) to mean “neighboring” elements of M : for instance $v_i \sim v_j$ means that v_i and v_j are neighboring vertices (share an edge), and similarly for pairs of faces, etc.

1.1 Basic Measures

Some geometric quantities can be easily and naturally measured directly from M . For instance (refer to figure 1, center, for labels):

- Edge lengths are just the Euclidean distance of their vertices, i.e. edge e_{01} has length $\|v_1 - v_0\|$.
- Face areas can also be computed from vertex positions: $\|(v_1 - v_0) \times (v_2 - v_0)\|$.
- Each triangle has a tangent plane spanned by any two of its edges, e.g. $\text{span}(v_1 - v_0, v_2 - v_0)$.
- Similarly each triangle has a normal vector $N = \frac{(v_1 - v_0) \times (v_2 - v_0)}{\|(v_1 - v_0) \times (v_2 - v_0)\|}$. Notice that the normal vector’s sign depends on the orientation of the face. The desire for consistently oriented normals motivates the requirement above that adjacent faces have consistent orientation.

But what about tangent planes or normal vectors at vertices and edges? Mean and Gaussian curvatures? We will approach defining these as we did for planar curves: we will take the list of properties of curvature from the smooth setting, discretize them, and use them as a basis for deriving discrete formulas which satisfy these properties *by construction*.

2 Discrete Functions

There are several natural discretizations of the space of functions over a surface. Here we will focus on discrete functions that assign a scalar to each vertex v_i of M : discrete functions can therefore be represented as vectors $F \in \Omega^0(M) \cong \mathbb{R}^n$, where $n = |V|$ is the number of vertices of M . As we did for functions on planar curves, we need a discrete inner product analogous to the smooth inner product $\langle f, g \rangle = \int_M fg \, dA$. The smooth inner product possesses several properties and symmetries:

- symmetry: $\langle f, g \rangle = \langle g, f \rangle$.
- bilinearity: $\langle f, \alpha g + h \rangle = \alpha \langle f, g \rangle + \langle f, h \rangle$.
- positivity: $\langle f, f \rangle \geq 0$, with equality only when $f = 0$ almost everywhere.
- locality: $\langle f, g \rangle$ depends only on f, g , and M in the same neighborhoods of the surface.
- $\langle 1, 1 \rangle = A(M)$, the total surface area of M .

Insisting on the first three properties for a discrete inner product gives us that this inner product must be of the form $\langle F, G \rangle = F^T A G$ for some positive-definite matrix A . Locality implies that A is diagonal, i.e.

$$\langle F, G \rangle = \sum_i F_i G_i A_i$$

for scalars A_i with $\sum_i A_i = A(M)$ by the last property. Two commonly-used possibilities for such vertex “area weight” A_i include:

- *barycentric area* $A_i = \sum_{f \sim v_i} \frac{1}{3} A(f)$: distribute each face’s surface area equally to its three boundary vertices. Clearly this partitions the surface’s total surface area, and (for nondegenerate faces) results in positive vertex area weights A_i .

- *circumcentric/Voronoi area*, where each triangle is divided into three quadrilaterals by drawing altitudes from the circumcenter to the triangle sides (see figure 1, right). These are the area weights that arise naturally from Discrete Exterior Calculus, which will be covered later in the course. However, they have complications that arise when triangles are obtuse: area weights can become negative (or difficult to compute, if “true” Voronoi areas are used to prevent this problem). On the other hand, circumcentric area weights are less sensitive to the surface meshing than barycentric areas; still barycentric areas are used most commonly in practice due to their robustness and simplicity.

3 Discrete Mean and Gaussian Curvature

With a notion of discrete functions and an inner product on discrete function space, we can proceed to discretize curvatures by invoking the properties of their smooth counterparts. Two properties in particular will be particularly useful: that the mean curvature normal is the gradient of surface area, and the Steiner expansion of volume enclosed by a surface as the surface is “inflated” in the normal direction.

3.1 Mean curvature from area

Recall that for a smooth, closed surface parameterized by $r(u, v)$, the gradient of surface area is given by

$$(\nabla_r A)(x) = 2H(x)N(x),$$

where H is the mean curvature and N the surface normal at a point x . Recall that this gradient is in the sense of the calculus of variations: for every (\mathbb{R}^3 -valued) variation δ over the surface,

$$\left. \frac{d}{dt} A(r + t\delta) \right|_{t \rightarrow 0} = \langle 2HN, \delta \rangle.$$

We will use this principle to discretize mean curvature H as a discrete function over a discrete surface M . Define $\nabla_v A$ to be the \mathbb{R}^3 -valued discrete function satisfying

$$\left. \frac{d}{dt} A(r + t\delta) \right|_{t \rightarrow 0} = \langle \nabla_v A, \delta \rangle \tag{1}$$

for all *discrete* variations δ . Here we extend the inner product on discrete scalar function to discrete vector-valued functions in the obvious way.

We want equation (1) to hold for *every* variation δ , so it must hold in particular for variations of one vertex v_i only:

$$\delta_j = \begin{cases} w, & i = j \\ 0, & i \neq j, \end{cases}$$

where $w \in \mathbb{R}^3$ is an arbitrary translation of vertex i . Plugging in this variation into equation (1) gives

$$\nabla_{v_i} \left(\sum_{f \sim v_i} A(f) \right) \cdot w = A_i (\nabla_v A)_i \cdot w$$

where the gradient on the left is the ordinary gradient in \mathbb{R}^3 , and most terms on both sides have vanished due to the sparsity of δ (on the right) and the fact that displacing v_i affects only the surface areas of the faces neighboring v_i (on the left). The above must hold for arbitrary w , which gives us a formula for $\nabla_v A$:

$$(\nabla_v A)_i = \frac{1}{A_i} \nabla_{v_i} \sum_{f \sim v_i} A(f).$$




Figure 2: *Left*: diagram for computing the triangle area gradient with respect to v_0 . *Right*: Figure illustrating the angles used in the “cotan formula” for edge e_{ij} .

To evaluate the gradient on the right, we need the triangle area gradient with respect to one vertex. Consider the diagram in figure 2. Letting A denote the area of this arbitrary triangle, we will compute $\nabla_{v_0} A$. First, we have that $A = \frac{1}{2}(b_1 + b_2)\|v_0 - p\|$, where p is the projection of v_0 onto the base. Moving v_0 parallel to the base does not change A , and moving v_0 out of plane does not affect the triangle area to first order, so we can assume the gradient $\nabla_{v_0} A$ is parallel to $v_0 - p$. Therefore

$$\nabla_{v_0} A = \frac{1}{2}(b_1 + b_2) \frac{v_0 - p}{\|v_0 - p\|}.$$

Substituting $p = \frac{b_1 v_2 + b_2 v_1}{b_1 + b_2}$ into the above gives

$$\begin{aligned} \nabla_{v_0} A &= \frac{b_1 v_0 + b_2 v_0 - b_1 v_2 - b_2 v_1}{2\|v_0 - p\|} \\ &= \frac{b_1}{2\|v_0 - p\|} (v_0 - v_2) + \frac{b_2}{2\|v_0 - p\|} (v_0 - v_1) \\ &= \frac{\cot \alpha}{2} (v_0 - v_2) + \frac{\cot \beta}{2} (v_0 - v_1). \end{aligned}$$

Therefore

$$2H_i N_i = (\nabla_v A)_i = \frac{1}{A_i} \sum_{v_j \sim v_i} \frac{\cot \alpha_{ij} + \cot \beta_{ij}}{2} (v_j - v_i), \quad (2)$$

where α_{ij} and β_{ij} are the triangle angles opposite edge e_{ij} (see figure 2, right).

The magnitude of the expression on the right of equation (2) gives us our discrete mean curvature H_i at v_i . As a bonus, except at vertices with zero mean curvature, its direction also gives us an expression for the surface normal N_i at v_i (up to sign; choosing the sign in practice is not difficult, and can be chosen so that e.g. its dot product with the average of the neighboring face normals is positive). N_i is often called the “area gradient” vertex normal, for obvious reasons.

The weights appearing above – the sum of the cotangents of the angles opposite edge e_{ij} are the infamous “cotan weights.” We will see them again when we study the discrete Laplace-Beltrami operator.

4 Curvatures from Inflation

Recall that for a smooth closed surface M that is sufficiently nice (e.g. is compact, injective, and an immersion), and for ϵ sufficiently small, the Steiner polynomial expansion of the volume enclosed by M is

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

where $M + \epsilon N$ is the surface N "inflated" by ϵ in the normal direction, $A(M)$ is the surface area of M , and H and K are the mean and Gaussian curvature, respectively. Notice that, by Gauss-Bonnet, the integral in the last term is a multiple of 4π and depends only on the genus of M .

Like we did for curves, once we decide what it means to inflate a closed mesh in the normal direction, we can impose the above formula and read off from it a definition of discrete mean and Gaussian curvature. For discrete curves, there were several ways of inflating: sweeping a disk of radius ϵ over the curve, moving all edges in their normal directions by ϵ , etc. The second approach will not work here: take a general mesh and look at the faces around a vertex. The vertex is the intersection of the planes containing each of the neighboring faces. If the vertex has three neighboring faces, i.e. valence three, then the three planes are guaranteed to meet at a vertex (barring certain degenerate corner cases) and it is not surprising that they do so. Four or more planes, however, generally do *not* meet at a vertex – they generally do not share a common point at all. So taking all neighboring faces of a vertex v and moving their planes by ϵ in their normal directions does not result in a well-defined inflated mesh, since the new faces do not meet at a vertex, except in special cases, like the platonic solids. (A mesh which *does* have the property that you can construct offset meshes with parallel faces is called a *conical mesh*, and these meshes are of particular interest in architectural geometry, where the two parallel faces might represent the front and back surfaces of a pane of glass, for instance.)

Instead, we can form an inflated surface by taking a ball B_ϵ of radius ϵ , and then taking the outer boundary of the Minkowski sum of the ball and M : in other words, sweeping the ball over M and then defining the outer boundary of the resulting solid as the new mesh $M + B_\epsilon$. What is the volume of this new mesh? To first order, we have

$$V(M + B_\epsilon) = V(M) + \epsilon A(M),$$

where $A(M)$ is the total area of the faces of M . This accounting ignores the volume contained in cylindrical pieces around each edge (of order ϵ^2 and in spherical caps around each vertex (of order ϵ^3). Using an argument identical to that used to calculate the area of the sectors of the inflated planar curve, the volume of the cylindrical pieces is

$$\sum_{e_{ij}} \frac{\psi_{ij}}{2\pi} \pi \epsilon^2 \|e_{ij}\| = \epsilon^2 \sum_{e_{ij}} \frac{\psi_{ij} \|e_{ij}\|}{2},$$

where ψ_{ij} is the *turning angle* of edge e_{ij} (the complement of the edge's dihedral angle). Notice that this sum naturally suggests a definition of mean curvature defined on mesh *edges* instead of vertices: we could build up a notion of discrete functions on edges, together with an inner product, and formalize this idea. We won't go into the details here, but this would lead to a *hinge-based* formula, linear in ψ_{ij} , for mean curvature, where mean curvature at each edge of the mesh depends only on that edge and the two neighboring triangles. This formulation is very popular in computer graphics, since it requires very little information about M to compute. But by being linear in ψ_{ij} , it asserts that one principal curvature direction is always in the direction of the mesh edge (with principal curvature zero), which is geometrically dubious – for this reason, although the hinge-based model of mean curvature is meaningful for calculating the *total* mean curvature of a surface (such as in the Steiner formula), it is not reliable for computing mean curvature *locally* near a particular edge. And indeed, for a sequence of discrete meshes converging to a smooth surface (even in a Sobolev sense, so that the mesh normals converge to the surface normals), the hinge-based discrete mean curvature is not guaranteed to converge to the smooth surface's mean curvature.

The third-order term in the volume expansion comes from spherical caps around each of the vertices. These caps are spherical polygonal sectors on a sphere of radius ϵ , where the spherical polygon's interior angles β_i (on the surface of the sphere) are the complements of the interior angles α_i of the triangles neighboring the vertex. To calculate the volume of the spherical cap, we begin by calculating the surface area of this spherical polygon: it turns out that this surface area is $\epsilon^2[(2 - n)\pi + \sum \beta_i]$, where n is the number of sides of the polygon. Notice that, unlike for Euclidean polygons, the area of the spherical polygon depends *only* on the angles of the polygon, and not on any lengths.

The surface area of the spherical cap is then $\epsilon^2(2\pi - \sum \alpha_i)$: the quantity $D_j = 2\pi - \sum \alpha_i$ is called the *angle deficit* and measures how far the neighborhood of a vertex v_j is away from being planar: the angle

deficit is zero for a flat neighborhood, positive for a conical neighborhood, and negative for a saddle-shaped neighborhood.

Finally, the volume ratio of the spherical cap to the total sphere is equal to the surface area ratio, and this gives us the volume of the spherical cap as

$$\frac{\epsilon^2 D_j}{4\pi\epsilon^2} \frac{4}{3}\pi\epsilon^3 = \frac{\epsilon^3}{3} D_j,$$

and the total volume contributed by all caps is

$$\frac{\epsilon^3}{3} \sum_{v_j} D_j.$$

We want this sum to equal the discrete analogue of total Gaussian curvature $\int_M K dA$, which we can write as the inner product $\langle K, 1 \rangle$. Since the discretization of this inner product is

$$\langle K, 1 \rangle = \sum_{v_j} K_j A_j,$$

equating terms in the two formulas yields the formula for Gaussian curvature, $K_i = \frac{D_i}{A_i}$. Note that since angle deficit is unitless, the discrete Gaussian curvature has units of inverse area, as expected.

Remark The above derivation assumed that the discrete mesh M is convex, since inflation fails to yield spherical caps at non-convex vertices. Unlike for the case of planar curves, this obstruction cannot be brushed aside, since there is no simple transformation (such as a reflection, in the case of the curve) which maps a discrete surface with a non-convex vertex to one where the vertex is convex. The angle-deficit formula is nevertheless commonly used even at non-convex vertices, and satisfies many of the properties of smooth Gaussian curvature – we will verify shortly that the angle-deficit formula for discrete Gaussian curvature has a discrete Gauss-Bonnet theorem. Still, it would be interesting to carefully examine the geometry at non-convex vertices, and study the properties of the resulting alternative formula for Gaussian curvature.

Remark Since the angle-deficit formula for K_i depends only on face angles and face area, and not on any extrinsic quantities such as dihedral angles, this discrete Gaussian curvature, like the smooth one, is *intrinsic* – it depends only on metric of the surface and not on the way it is embedded. Moreover, the angle-deficit Gaussian curvature satisfies a discrete analogue of Gauss-Bonnet: to see this, note that for a closed mesh M

$$\langle K_i, 1 \rangle = \sum_{v_i} (2\pi - \sum \alpha_j) = 2\pi|V| - \pi|F|,$$

where $|V|$ and $|F|$ are the number of vertices and faces in M , respectively: the sum counts every angle of every face exactly once, so the total of all of the angles α must be pi (the sum of the interior angles of one triangular face) times the number of faces. Next, for a triangular mesh, we can associate to each face three neighboring “half-edges” that it shares with one of its neighboring faces; therefore $3|F| = 2|E|$ and

$$\langle K_i, 1 \rangle = 2\pi|V| - \pi(|F| + 2|E| - 3|F|) = 2\pi(|V| - |E| + |F|).$$

But $|V| - |E| + |F|$ is exactly the Euler characteristic $2 - 2g$ of the mesh, and so

$$\langle K_i, 1 \rangle = 4\pi(1 - g),$$

in exact accord with the smooth Gauss-Bonnet theorem.