Chapter 4
Computations on Triangulated Surfaces

4.1 Triangulated Surfaces

4.1.1 Definition and Notation

Triangulated surfaces provide a three-dimensional generalization of polygons in two dimensions. Surfaces are usually stored on computers in this form, and these are the kinds of objects that must be handled in practical applications.

In full generality, a triangulated surface is a set of *vertices* $\mathcal{V} = \{v_1, \dots, v_M\}$ with a family of 3-tuples of indices $\mathcal{F} = \{f_1, \dots, f_K\}$, where each f_k takes the form $f_k = (j_{k1}, j_{k2}, j_{k3}) \in \{1, \dots, M\}$. One associates to f_k the triangle (or face) in the triangulation defined by $F_k = (v_{k1}, v_{k2}, v_{k3})$, using the abbreviated notation $v_{kl} := v_{j_{kl}}$. The set of edges of the triangulation is the family of unordered pairs of vertices which belong to the same face and will be denoted by $\mathcal{E} = \{e_1, \dots, e_Q\}$.

The order of the vertices in each face is important and defines its orientation, which is invariant up to a cyclic permutation of the vertices. We will only consider *regular* triangulations, which are such that the intersection of two faces is either empty or an edge. This excludes those situations in which the contact between two faces occurs at a vertex only, or in which some vertex belongs to the interior of an edge. The number

$$\chi = |\mathcal{V}| - |\mathcal{E}| + |\mathcal{F}|$$

is a topological invariant of the surface called the Euler characteristic.

For a vertex v_i, we let $\mathcal{F}_i$ denote the set of indexes of faces that contain it, and $\mathcal{E}_i$ the set of indexes of edges that contain it. We also let $\mathcal{V}_i$ denote the set of indexes of vertices (distinct from v_i) that belong to one of the 3-tuples in $\mathcal{F}_i$. $(\mathcal{V}_i, \mathcal{E}_i, \mathcal{F}_i)$ represents the neighborhood of v_i in the triangulation.

L. Younes, *Shapes and Diffeomorphisms*, Applied Mathematical Sciences 171,
https://doi.org/10.1007/978-3-662-58496-5_4

The triangulation is said to be consistent if, whenever two faces intersect, their common edge is ordered in different directions in the two faces. A consistent triangulation is the equivalent of an oriented surface. We only consider consistent triangulations in the following.

4.2 Estimating the Curvatures

Given a triangulated surface, the next step is to compute differential descriptors, and in particular discrete forms of the curvatures. We address this problem in this section, focusing on a few important methods that have recently emerged in the literature.

4.2.1 Taylor Expansions

The unit normal to an oriented triangle (v_1, v_2, v_3) is the vector

$$N = \frac{(v_2 - v_1) \times (v_3 - v_1)}{|(v_2 - v_1) \times (v_3 - v_1)|}.$$

Each face F_k in the triangulation therefore carries a uniquely defined normal, N_k^f. We can associate the normal to a specific point inside the face, for example its centroid $(v_1 + v_2 + v_3)/3$. (There are several possible definitions of the center of a triangle, however, including the circumcenter, which is the center of the circumscribed circle, the incenter, which is the center of the inscribed center, or the orthocenter, the intersection of the lines passing through the vertices and orthogonal to the opposite edge.)

In many cases, one also wants to define normals at the vertices. This can be done using a weighted average of the normals at the neighboring faces. If v_i is a vertex, define

$$N_i^v = \frac{\sum_{k \in \mathcal{F}_i} w_i(F_k) N_k^f}{|\sum_{k \in \mathcal{F}_i} w_i(F_k) N_k^f|},$$

where $w_i(F_k)$ gives a measure of the "importance" of face F_k relative to vertex v_i. The simplest definition is the area, area(F_k), independent in this case of the chosen vertex. In [194], it is suggested to use the area of the part of the face which is closer to v_i than to any of the other two vertices. This is the intersection of the face F_k with the region delimited by the following four points: v_i, the two midpoints of the edges of F_k that contain v_i and the circumcenter of F_k. Notice that the circumcenter lies outside of F_k if the triangle is obtuse, as illustrated in Fig. 4.1. Such regions form Voronoï cells. Let F_{ki} denote the part of face F_k which is associated to v_i in this way. One can use $w_i(F_k) = \text{area}(F_{ki})$.

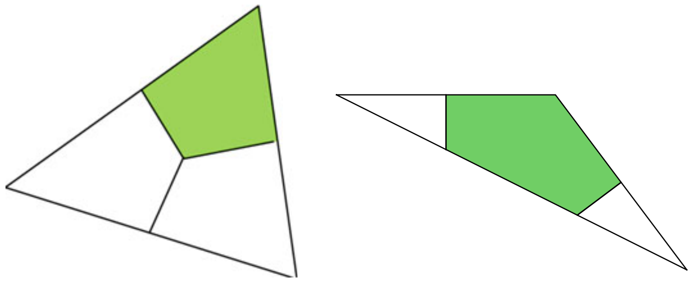

Fig. 4.1 Decomposition of triangles into Voronoï cells when the circumcenter is interior to the triangle (left) and when it is exterior (right)

Similarly, we can define a normal along an edge e to be a weighted average of the normals to the faces that intersect at e, using, for example, the areas of the faces as weights.

Having an estimation of the normal at each vertex allows for the approximation of the normal curvature of a curve on the surface passing through this vertex, which yields the second fundamental form. If $j \in \mathcal{V}_i$, the two-point path (v_i, v_j) provides a discrete curve fragment passing through v_i. Define the tangent vector $T_{ij} = (v_j - v_i)/|v_j - v_i|$. Using Definition 3.26, one possible approximation of the second fundamental form at the midpoint between v_i and v_j in the direction T_{ij} (which is also the normal curvature of the curve fragment at the midpoint) is

$$II_{ij} := -T_{ij}^T \left(\frac{N_j^v - N_i^v}{|v_j - v_i|} \right) = -\frac{(N_j^v - N_i^v)^T (v_j - v_i)}{|v_j - v_i|^2}.$$

Also, using a Taylor expansion (assuming that N^v is the restriction to the vertices of a smooth function), one can prove (the justification being left to the reader) that

$$\frac{(N_j^v + N_i^v)^T (v_j - v_i)}{|v_j - v_i|^2} = O(|v_j - v_i|),$$

and adding this expression to the previous estimate II_{ij} yields the alternative formula [273]

$$II_{ij} := \frac{2(N_i^v)^T (v_j - v_i)}{|v_j - v_i|^2}.$$

Because the matrix dN is symmetric in the tangent plane, it is described by three parameters in any orthonormal basis. Since each computation of the discrete second fundamental form yields one linear equation involving dN, this requires at least three edges for its estimation, which is the minimum number provided by the triangulation. One possible way to estimate curvatures is to select an arbitrary basis (a_i, b_i) of the tangent plane to the surface at v_i, $T_{v_i}M$, which is by definition the plane perpendicular to N_i^v (for example, assuming that N_i^v is not parallel to the x-axis, take $a_i = (1, 0, 0)^T \times N_i^v$ and $b_i = N_i^v \times a_i$); then, compute, for each $j \in \mathcal{V}_i$, the coordinates (x_{ij}, y_{ij}) of the normalized orthogonal projection of T_{ij} onto this basis. We have

$$x_{ij} = a_i^T T_{ij} / \sqrt{(a_i^T T_{ij})^2 + (b_i^T T_{ij})^2}$$

$$\text{and } y_{ij} = b_i^T T_{ij} / \sqrt{(a_i^T T_{ij})^2 + (b_i^T T_{ij})^2}.$$

Then, letting $dN_i = \begin{pmatrix} \alpha_i & \gamma_i \\ \gamma_i & \beta_i \end{pmatrix}$ in this basis, we have the system of linear equations

$$\alpha_i x_{ij}^2 + 2\gamma_i x_{ij} y_{ij} + \beta_i y_{ij}^2 = -II_{ij}, \; j \in \mathcal{V}_i.$$

This is an over-constrained system, for which one can compute a least-squares solution. Once dN_i is computed, its trace, determinant and eigenvalues provide an estimation of the mean, Gaussian and principal curvatures.

A more direct approach to estimating the curvature from the second fundamental form has been proposed in [273]. Introduce, for continuous surfaces, the matrix (defined at a point p in the surface)

$$\Sigma_p = \frac{1}{2\pi} \int_0^{2\pi} \kappa_N(T_\theta) T_\theta T_\theta^T \, d\theta,$$

where T_θ is the rotation (within the tangent plane) of an arbitrary reference vector $T \in T_p M$ by an angle θ. A direct computation of this integral (using the basis $(T, T_{\pi/2})$) shows that

$$\Sigma_p = \frac{3}{8} dN(p) - \frac{1}{8} \det(dN(p)) dN(p)^{-1},$$

the last term being the adjugate matrix of $dN(p)$ (therefore also defined when $dN(p)$ is singular).

This implies that the eigenvalues of Σ_p are $\lambda_1 = -(3\kappa_1(p) - \kappa_2(p))/8$ and $\lambda_2 = -(3\kappa_2(p) - \kappa_1(p))/8$ (which can be used to compute the curvatures), and that the eigenvectors of Σ_p coincide with those of $dN(p)$ and therefore provide the principal directions.

Returning to the discrete case, the curvatures at vertex v_i can therefore be estimated from an approximation Σ_i of Σ_{v_i}. Such an approximation is provided by the simple formula

$$\Sigma_i = -\sum_{j \in \mathcal{V}_i} w_{ij} II_{ij} \, T_{ij} T_{ij}^T \Big/ \sum_{j \in \mathcal{V}_i} w_{ij},$$

where $w_{ij} = (w_i(F_j^+) + w_i(F_j^-))/2$, F_j^+ and F_j^- being the faces that contain the edge $\{v_i, v_j\}$.

4.2.2 Gauss–Bonnet and Area Minimization

In the previous section, the curvature computations were based on Taylor expansions of formulas that apply on smooth surfaces. More recently [85], an increased focus has been made on obtaining expressions that derive from intrinsic properties of surfaces that can be extended to polyhedral surfaces.

The right-hand side of Eq. (3.32) in the Gauss–Bonnnet theorem can still be defined on polyhedral surfaces. This fact is used in [194] to provide an approximation of the Gauss curvature, using, for a vertex v_i in the triangulation, the region A_i formed by the union of the Voronoï cells around v_i (Fig. 4.1). The expression is very simple because, in both cases in Fig. 4.1, the sum of the (one or two) exterior angles in each part of A_i coincides with the angle of the corresponding face at v_i. For $k \in G_i$, denoting by θ_{ik} the angle of the face F_k at v_i, we see that the right-hand side of (3.32) is given by $2\pi - \sum_{k\in\mathcal{F}_i} \theta_{ik}$. Approximating K by a constant over A_i, we get the formula

$$K_i = \frac{1}{|A_i|}\Big(2\pi - \sum_{k\in\mathcal{F}_i} \theta_{ik}\Big).$$

The area, $|A_i|$, can be computed in closed form. It is the sum of the areas of the shaded regions in Fig. 4.1, over all faces that contain v_i. Let as above θ_{ik} be the angle at v_i for a face F_k. Let v'_{ik} and v''_{ik} be the other two vertices of F_k so that v_i, v'_{ik} and v''_{ik} are ordered consistently with the orientation of F_k. Let e'_{ik} be the edge opposite v'_{ik} in F_k and e''_{ik} the edge opposite v''_{ik} (we will later denote by e_{ik} the edge opposite v_i). Finally, let θ'_{ik} and θ''_{ik} respectively denote the angles at v'_{ik} and v''_{ik}. Then, the area, a_{ik}, of the shaded region in Fig. 4.1 in the acute case is given by

$$a_{ik} = \frac{1}{8}(|e'_{ik}|^2 \operatorname{ctn}(\theta'_{ik}) + |e''_{ik}|^2 \operatorname{ctn}(\theta''_{ik})).$$

In the obtuse case, and if θ_{ik} is the obtuse angle,

$$a_{ik} = \frac{1}{2}|e'_{ik}|\,|e''_{ik}| \cos\theta_{ik} - \frac{1}{8}(|e'_{ik}|^2 \cos\theta''_{ik} + |e''_{ik}|^2 \cos\theta'_{ik}).$$

Finally, still in the obtuse case, and when θ_{ik} is one of the acute angles,

$$\frac{1}{8}|\tilde{e}_{ik}|^2 \cos\theta_{ik}$$

where $\tilde{e}_{ik}$ is the side opposed to the other acute vertex.

Given this, $|A_i|$ is the sum of these areas over $k \in G_i$. When there is no obtuse triangle around v_i, the area $|A_i|$ has another simple expression [194]. For $l \in \mathcal{E}_i$ (edges stemming from v_i), let α_{il} and β_{il} be the angles at vertices opposed to e_l in the triangles that intersect at e_l. Then

$$|A_i| = \frac{1}{8} \sum_{l \in \mathcal{E}_i} (\text{ctn}\,\alpha_{il} + \text{ctn}\,\beta_{il})|e_l|^2.$$

To address the mean curvature, we first use an important interpretation of it as a "gradient" of the surface area. Let S be a surface and $h : S \to \mathbb{R}^3$ be a (smooth) vector field on S. Assume that $h = 0$ on the boundary of S (if S has one). Define the surface S_ε as the one obtained by displacing each $p \in S$ along the vector $\varepsilon h(p)$. Then (this will be proved in Proposition 5.4)

$$\partial\,(\text{area}(S_\varepsilon))_{|\varepsilon=0} = 2 \int_S H(p)\, h(p)^T N(p)\, d\sigma_S(p).$$

One can make the same construction with a discrete surface Σ by associating to each vertex v_i a small displacement $\varepsilon h_i \in \mathbb{R}^3$, and computing the derivative of the area of the obtained surface Σ_ε. Approximating the right-hand side in the above formula, we will then identify:

$$\partial\,(\text{area}(\Sigma_\varepsilon))_{|\varepsilon=0} = 2 \sum_{i=1}^{M} h_i^T (H_i N_i)|A_i|, \tag{4.1}$$

where A_i is the neighborhood attributed to v_i and $H_i N_i$ can then be interpreted as the discretized product of the mean curvature with the normal at v_i.

Given that the area of a triangle with vertices v_1, v_2, v_3 is given by the half-norm of the cross product $(v_2 - v_1) \times (v_3 - v_1)$, the left-hand side in (4.1) is

$$\begin{aligned}
&\frac{1}{2} \sum_{k=1}^{K} ((h_{2k} - h_{1k}) \times (v_{k3} - v_{k1}) + (v_{k2} - v_{k1}) \times (h_{k3} - h_{k1}))^T N_k^f \\
&= \frac{1}{2} \sum_{k=1}^{K} (h_{k1} \times (v_{k2} - v_{k3}) + h_{k2} \times (v_{k3} - v_{k1}) + h_{k3} \times (v_{k2} - v_{k1}))^T N_k^f,
\end{aligned}$$

where h_{k1}, h_{k2} and h_{k3} are the displacements associated with the vertices of F_k and N_k^f is the normal to F_k.

For $k \in \mathcal{F}_i$, let e_{ik} be the oriented edge opposite v_i. Using the relation $(x \times y)^T z = x^T (y \times z)$ and reordering the sums, we can write

$$\frac{d}{d\varepsilon} \text{area}(\Sigma_\varepsilon)_{|\varepsilon=0} = \frac{1}{2} \sum_{i=1}^{M} h_i^T \left(\sum_{k \in \mathcal{F}_i} e_{ik} \times N_k^f \right).$$

This provides a definition of the discrete mean curvature at v_i:

$$H_i N_i = \frac{1}{4|A_i|} \sum_{k \in \mathcal{F}_i} e_{ik} \times N_k^f.$$

Reordering this sum over edges and explicitly computing the cross product leads to the equivalent expression [194]

$$H_i N_i = \frac{1}{4|A_i|} \sum_{l \in \mathcal{E}_i} (\operatorname{ctn} \alpha_{il} + \operatorname{ctn} \beta_{il}) e_l,$$

where α_{il} and β_{il} are, as before, the angles at the vertices opposite to e_l in each of the faces that contain e_l (e_l being oriented from v_i to the other vertex).

Note that this computation provides an estimate of the normal and the mean curvature together.

4.2.3 Curvature Measures

The Smooth Case

There is another way to interpret curvature on a surface that can be generalized to the non-smooth case, leading to another formula for curvature approximation on triangulated surfaces. On smooth surfaces, this is related to the volume of so-called parallel sets. We first show that smooth surfaces have positive reach, in a discussion that parallels the one in Sect. 1.13.2. We use the same notation as in that section, letting, for a surface M,

$$d_M(p) = \operatorname{dist}(p, M) = \inf \{|p - q| : q \in M\},$$

$\mathcal{U}_M$ be the set of points p that have a unique closest point, $\pi_M(p)$, on M, $r(M, q)$ be the supremum of the radii of balls centered at q included in $\mathcal{U}_M$ and $r(M)$ be their minimum over $q \in M$ (the reach of M). Propositions 1.20, 1.21 and 1.22 remain true in the present case, as does the fact that d_M is differentiable on $\mathring{\mathcal{U}}_M \setminus M$. We prove a version of Proposition 1.23 for surfaces.

Proposition 4.1 *Let M be a closed C^2 regular surface. Then, we have the following statements.*

(i) If $|p - q| = d_M(p)$ $(q \in M)$ then $p = q + tN_M(q)$ with $|t| = d_M(p)$ and $\max(t\kappa_1(q), t\kappa_2(q)) \leq 1$, where κ_1 and κ_2 are the principal curvatures.

(ii) Let

$$\rho_M = \max \left\{ \frac{2\,|(\tilde{q} - q)^T N_M(q)|}{|\tilde{q} - q|^2} : q, \tilde{q} \in M, q \neq \tilde{q} \right\}. \qquad (4.2)$$

Then $\rho_M < \infty$ and $r(M) \geq 1/\rho_M > 0$. In particular, $\mathring{\mathcal{U}}_M$ is not empty.

(iii) The distance map is differentiable on $\mathring{\mathcal{U}}_M \setminus M$.

Proof The proof is similar to that of Proposition 1.23 and we only highlight the differences. To prove (i), take an arc-length parametrized curve γ on M such that $\gamma(0) = q$

and let $f(t) = |p - \gamma(t)|^2$. Then, $\dot{f}(0) = -2(p-q)^T\dot{\gamma}(0)$, which much vanish for all γ, showing that $p - q$ is perpendicular to T_qM and therefore proportional to $N_M(q)$, i.e., $p = q + tN_M(q)$ for some t. Taking the second derivative of f (which must be non-negative at 0) yields $\ddot{f}(0) = 2 - 2(p-q)^T\ddot{\gamma}(0) = 2(1 - t\kappa_N^{(\gamma)}(0))$, so that $t\kappa_N^{(\gamma)}(0) \le 1$. It then suffices to use the fact that $\kappa_1(q) \le \kappa_N^{(\gamma)}(0) \le \kappa_2(q)$.

To prove that ρ_M is finite, take $q_n, \tilde{q}_n$ such that

$$c_n := \frac{2\,|(\tilde{q}_n - q_n)^T N_M(q_n)|}{|\tilde{q}_n - q_n|^2}$$

tends to infinity, which implies that subsequences can be taken such that $q_n, \tilde{q}_n \to q$. Take a local chart around q such that $q_n = m(u_n, v_n)$ and $\tilde{q}_n = m(\tilde{u}_n, \tilde{v}_n)$. Let γ_n be a minimizing geodesic such that $\gamma_n(0) = q_n$ and $\gamma(s_n) = \tilde{q}_n$. Then the same argument as that of Proposition 1.23 can be used to prove that $|\tilde{q}_n - q_n|^2 = s_n^2 + o(s_n^2)$ and

$$|(\tilde{q}_n - q_n)^T N_M(q_n)| = |\kappa_N^{(\gamma_n)}(0)|s_n^2 + o(s_n^2),$$

contradicting the assumption that $c_n \to \infty$. The rest of the proof of (ii) is identical to Proposition 1.23.

For (iii), one shows that, if $p \in \mathring{\mathcal{U}}_M \setminus M$ and $q = \pi_M(p)$, then $1 - \max(t\kappa_1(q), t\kappa_2(q)) > 0$ with the same argument as in Proposition 1.23. Take a positively oriented chart $(u, v) = m(u, v)$ around q and consider the mapping $\varphi_m(u, v, t) = m(u, v) + tN_M(m(u, v))$. Then, letting $N_m = N_M \circ m$,

$$\partial_1\varphi_m = \partial_1 m + t\partial_1 N_m,\ \partial_2\varphi_m = \partial_2 m + t\partial_2 N_m,\ \partial_3\varphi_m = N_m$$

so that

$$\begin{aligned} \det(d\varphi_m) &= (\partial_1\varphi_m \times \partial_2\varphi_m)^T \partial_3\varphi_m \\ &= (\partial_1 m \times \partial_2 m)^T N_m + t(\partial_1 m \times \partial_2 N_m + \partial_1 N_m \times \partial_2 m)^T N_m \\ &\quad + t^2(\partial_1 N_m \times \partial_2 N_m)^T N_m. \end{aligned} \tag{4.3}$$

We have $\partial_1 m \times \partial_2 m = |\partial_1 m \times \partial_2 m|\, N_m$. Moreover, for any linear operator A on $\mathbb{R}^3$ and any basis (u_1, u_2, u_3) in $\mathbb{R}^3$, we have (the proof being left to the reader)

$$(u_1 \times u_2)^T A u_3 + (u_2 \times u_3)^T A u_1 + (u_3 \times u_1)^T A u_2 = \det(u_1, u_2, u_3)\mathrm{trace}(A).$$

Applying this to

$$(\partial_1 m \times \partial_2 N_m + \partial_1 N \times \partial_2 m)^T N_m = (N_m \times \partial_1 m)^T \partial_2 N_m + (\partial_2 m \times N_m)^T \partial_1 N_m$$

with $\partial_1 N_m = dN_m\, \partial_1 m$, $\partial_2 N_m = dN_m\, \partial_2 m$, taking $A = dN_m$ on T_pM and $AN_m = 0$, we get

$$(\partial_1 m \times \partial_2 N_m + \partial_1 N_m \times \partial_2 m)^T N_m = |\partial_1 m \times \partial_2 m| \operatorname{trace}(dN_m)$$
$$= -2|\partial_1 m \times \partial_2 m|\, H(m),$$

where H is the mean curvature. Moreover, since $\partial_1 N_m$ and $\partial_2 N_m$ are tangent to M at m, $(\partial_1 N_m \times \partial_2 N_m)^T N_m$ is the two-dimensional determinant of $[dN_m\, \partial_1 m, dN_m\, \partial_2 m]$, therefore equal to $K(m)|\partial_1 m \times \partial_2 m|$. We therefore have

$$\begin{aligned}\det(d\varphi_m) &= (1 - 2tH(m) + t^2 K(m))|\partial_1 m \times \partial_2 m| \\ &= (1 - t\kappa_1(m))(1 - t\kappa_2(m))|\partial_1 m \times \partial_2 m|. \end{aligned} \quad (4.4)$$

The determinant is therefore positive, and the differentiability of d_M at p can be obtained using the inverse function theorem, as done in the proof of Proposition 1.23 □

Proposition 4.1 ensures that the mapping

$$\begin{aligned}\varphi_M: \quad M \times (-r, r) &\to \mathbb{R}^3 \\ (q, t) &\mapsto q + tN_M(q)\end{aligned}$$

is one-to-one for $r < r(M)$ onto $V_M(r) = \{p : d_M(p) < r\}$. More generally, for $B \subset M$, consider the sets $V_r(M, B) = \varphi_M(B \times (-r, r))$ and $V_r^+(M, B) = \varphi_M(B \times (0, r))$. Using the fact that φ_m introduced in the proof of Proposition 4.1 is such that $\varphi_m(u, v, t) = \varphi_M(m(u, v), t)$ and assuming that M is entirely covered by a local chart, Eq. 4.4 implies that

$$\begin{aligned}&\mathrm{Vol}(V_r^+(M, B)) \\ &= \int_0^r \int_{m^{-1}(B)} (1 - 2tH(m(u, v)) + t^2 K(m(u, v)))|\partial_1 m \times \partial_2 m|\, du\, dv \\ &= r\,\mathrm{Area}(B) - r^2 \int_B H\, d\sigma_M + \frac{r^3}{3} \int_B K d\sigma_M, \end{aligned} \quad (4.5)$$

where σ_M is the volume measure on M. The last expression for $\mathrm{Vol}(V_r^+(M, B))$ remains true even when B is not completely covered by a local chart, as can easily be proved by using partitions of unity, or covering B by a union of local charts up to a set of measure zero.

This discussion can also be extended to compact surfaces with boundary. Similar to the case of curves, we need to generalize the definition of normal vectors to boundary points. Let $p \in \partial M$ and $N_{\partial M}(p) \in T_p M$ denote the unit normal to the boundary at p pointing inward (toward M). Then, a vector N is normal to M at p if it can be written in the form

$$N = t_1 N_{\partial M}(p) + t_2 N_M(p)$$

with $t_1 \leq 0$. We let $\mathcal{N}_M(p)$ denote the set of unit normals at p. Then the statements of Proposition 4.1 remain true, provided that (4.2) is replaced by

$$\rho_M = \max\left\{\frac{2\,(\tilde{q}-q)^T N}{|\tilde{q}-q|^2} : q, \tilde{q} \in M, q \neq \tilde{q}, N \in \mathcal{N}_M(q)\right\}.$$

Equation (4.5) remains true whenever $B \subset M$. For $B \subset \partial M$, the computation must be modified. Represent $B \subset \partial M$ as a parametrized curve $\gamma : (0, L) \to \mathbb{R}^3$ and define

$$\varphi(s, t_1, t_2) = \gamma(s) + t_1 N_{\partial M}(\gamma(s)) + t_2 N_M(\gamma(s)).$$

Assume that γ is parametrized by arc length and oriented so that $\dot{\gamma} \times N_{\partial M} = N_M$. Consider the set $V_r(M, B) = \varphi(B \times \Gamma_r)$ where Γ_r is the half disc $\{(t_1, t_2) : t_1^2 + t_2^2 < r^2, t_1 \leq 0\}$, so that $V_r(M, B)$ is the set of points in $\mathbb{R}^3$ that have closest point to M in B at distance less than r. We have

$$\begin{aligned}\det(d\varphi(s, t_1, t_2)) &= \det(\dot{\gamma}(s), N_{\partial M}(\gamma(s)), N_M(\gamma(s)))\\ &\quad + t_1 \det(\partial_s N_{\partial M}, N_{\partial M}(\gamma(s)), N_M(\gamma(s)))\\ &\quad + t_2 \det(\partial_s N_M, N_{\partial M}(\gamma(s)), N_M(\gamma(s)))\\ &= 1 - t_1 \kappa_g^{(\gamma)}(s) - t_2 \kappa_N^{(\gamma)}(s).\end{aligned}$$

Indeed, we have $N_{\partial M} \times N_M = \dot{\gamma}$ and

$$\det(\partial_s N_{\partial M}, N_{\partial M}(\gamma(s)), N_M(\gamma(s))) = \partial_s N_{\partial M}^T \dot{\gamma} = -N_{\partial M}^T \ddot{\gamma} = -\kappa_g^{(\gamma)}.$$

Similarly

$$\det(\partial_s N_M, N_{\partial M}(\gamma(s)), N_M(\gamma(s))) = \partial_s N_M^T \dot{\gamma} = -\kappa_N^{(\gamma)}.$$

We can now compute (introducing radial coordinates)

$$\begin{aligned}\mathrm{Vol}(V_r(M, B)) &= \int_0^L \int_{t_1^2+t_2^2<r^2, t_1<0} (1 - t_1 \kappa_g^{(\gamma)}(s) - t_2 \kappa_N^{(\gamma)}(s))\, d\theta\, d\rho\, ds\\ &= \int_0^L \int_0^r \int_{-\pi/2}^{\pi/2} (1 + \rho\cos\theta \kappa_g^{(\gamma)}(s) - \rho \sin\theta \kappa_N^{(\gamma)}(s))\, d\theta\, d\rho\, ds\\ &= r\,\mathrm{length}(B) + r^2 \int_0^L \kappa^{(\gamma)}(s)\, ds.\end{aligned}$$

Applied to $r < r(M)$, Eq. (4.5) provides a new interpretation of the integrals of mean and Gauss curvatures over subsets of M in terms of the volumes of the sets $V_r^+(M, B)$ (which are often called *parallel sets* along the surface M). These parallel sets may be defined for sets that are much more general than smooth surfaces and their volumes then lead to generalized versions of the curvature. We now see how these ideas can be applied to triangulated surfaces.

The Discrete Case

We first generalize the definition of $V_B(r)$ to the non-smooth case, for which we will need a generalized definition of the set of unit normals to M at a given point. This can be done in two equivalent ways. Here M is an arbitrary closed set with positive reach, i.e., such that $r(M) > 0$.

The first point of view is to let, for $r < r(M)$,

$$V_r(M, B) = \{p : 0 < d_M(p) < r \text{ and } \pi_M(p) \in B\}.$$

When M is the boundary of a compact set Ω (e.g., when M is a closed surface), we can define

$$V_r^+(M, B) = V_r(\Omega, B),$$

still for $B \subset M$.

Federer [106] has proved that (4.5) can be generalized to sets of positive reach, so that $\mathrm{Vol}(V_B^+(r))$ is polynomial in r, taking the form

$$\mathrm{Vol}(V_r^+(M, B)) = r\mu_0(M, B) - r^2\mu_1(M, B) + \frac{r^3}{3}\mu_2(M, B). \tag{4.6}$$

In particular, $\mu_1(M, B)$ and $\mu_2(M, B)$ are generalizations of the integrals of the curvatures on B, and are called the mean and Gauss *curvature measures* on M. They have the important property of being additive, satisfying in particular

$$\mu_i(M \cup M', B) = \mu_i(M, B) + \mu_i(M', B) - \mu_i(M \cap M', B). \tag{4.7}$$

Although it already is a rich class of sets (including, for example, all convex sets), sets of positive reach do not include non-convex polyhedrons, so the construction cannot be immediately extended to triangulated surfaces. But formula (4.7) provides the key for this extension. Indeed, one can define a union of sets of positive reach as a set M that can be decomposed into

$$M = \bigcup_{j \in J} M_j,$$

where each M_j has positive reach and any nonempty intersection of M_j's has positive reach [314]. Then, iterating (4.7) (using the inclusion-exclusion formula), we can set

$$\mu_k(M, B) = \sum_{I \subset J} (-1)^{|I|-1} \mu_k\Big(\bigcap_{j \in I} M_j, B\Big), \tag{4.8}$$

the left-hand side being well-defined from the hypotheses. This is a valid definition of the right-hand side because it can be shown that the result does not depend on the chosen decomposition of M, which is not unique. This extension now includes all polyhedrons (and interiors of compact triangulated surfaces).

The second point of view gives an alternative interpretation of the curvature measures, based on the *normal bundle* to a set M. This requires a general definition of tangent and normal vectors to an arbitrary set $M \subset \mathbb{R}^3$ [107]. We already gave a general definition of tangent vectors in Definition 1.5, which defined the tangent set to M at $p \in M$ as the set T_pM of vectors $v \in \mathbb{R}^d$ such that, for any $\varepsilon > 0$, there exist $x \in M$ and $r > 0$ such that $|x - p| < \varepsilon$ and $|v - r(x - p)| < \varepsilon$.

We will later use the fact that, when M is included in the boundary of an open set Ω, $T_m\Omega$ for $m \in M$ contains vectors in T_mM and all vectors v such that $m + \varepsilon v \in \Omega$ for small ε (vectors that point to the interior of Ω). The special cases that follow will be important when studying triangulated surfaces.

Single points: Assume that $M = \{a\}$. It is clear from the definition that any tangent vector to M must satisfy $|v| < \varepsilon$ for any $\varepsilon > 0$ so that $T_aM = \{0\}$.

End-points of curves: Let $\gamma : [0, 1] \to \mathbb{R}^3$ be a smooth regular curve, $M = \gamma([0, 1])$ and $a = \gamma(0)$. Then, any $x \in M$ close to a is equal to $\gamma(u)$ for $u \simeq 0$, and a tangent vector v at a must be such that $v \simeq r(\gamma(u) - \gamma(0))$ with $r > 0$, so that T_aM is the half-line $\mathbb{R}^+\dot\gamma(0)$. By the same argument, if $b = \gamma(1)$, $T_bM = \mathbb{R}^-\dot\gamma(1)$. (Of course, the tangent space at interior points is the full line generated by the tangent vector $\dot\gamma$.)

Triangles: If M a triangle (including its interior) and a is on its boundary, then T_aM is simply the set of vectors v such that $a + v$ points towards the interior of M.

Normals can now be derived from tangents, as stated in the next definition.

Definition 4.2 Let $M \subset \mathbb{R}^3$. For $p \in M$, the normal vectors to M at p form the set $\mathcal{N}_pM$, containing all vectors n such that $n^Tv \le 0$ for $v \in T_pM$.

The normal bundle of M is the set $\mathcal{N}M \subset \mathbb{R}^3 \times \mathbb{R}^3$ defined by

$$\mathcal{N}M = \left\{(p, n), p \in M, n \in \mathcal{N}_pM, |n| = 1\right\}.$$

When $M \subset \partial\Omega$, we can also consider

$$\mathcal{N}M^+ = \left\{(p, n), p \in M, n \in \mathcal{N}_p\Omega, |n| = 1\right\}.$$

This corresponds to normals to M pointing outward from Ω.

The normal bundle is the structure on which the new curvature measures will be defined. Let us describe it for the previous examples. First, if M is a smooth closed oriented surface in $\mathbb{R}^3$, $\mathcal{N}M$ is simply the set $\{(p, N(p)), p \in M\} \cup \{(p, -N(p)), p \in M\}$. The set $\mathcal{N}M^+$ only contains elements $(p, -N(p))$ (assuming that M is positively oriented).

If M is a closed curve, with regular parametrization $s \mapsto \gamma(s)$, then $\mathcal{N}M = \left\{(\gamma(s), n) : n^T\dot\gamma = 0, |n| = 1\right\}$ (this can be thought of as a tube centered around γ).

If $M = \{a\}$, then (since $T_aM = \{0\}$), $\mathcal{N}M = \{a\} \times S^2$, where S^2 is the unit two-dimensional sphere.

When M is an open curve, parametrized by γ, the set $\mathcal{N}_aM$ for $a = \gamma(0)$ is the half-sphere $S^2 \cap \{n^T\dot\gamma(0) \le 0\}$, while, for $b = \gamma(1)$, it is $\mathcal{N}_bM = S^2 \cap \{n^T\dot\gamma(1) \ge 0\}$. The whole set $\mathcal{N}M$ can be thought of as a "sausage" around γ.

Finally, if M is a triangle, and a is on an edge, but not at a vertex, $\mathcal{N}_a M$ is a half-circle, being the intersection of the unit circle orthogonal to the edge and the half-space $\{n^T \nu \geq 0\}$, where ν is normal to the edge, in the triangle plane, and pointing towards the interior of the triangle. If a is a vertex, the set $\mathcal{N}_a M$ is the "spherical corner" formed by the intersection of the unit sphere and the two half-planes $\{n^T e \leq 0\}$ and $\{n^T e' \leq 0\}$, where e and e' are the two edges stemming from a (oriented from a to the other vertex).

The interesting fact in the previous examples is that, in each case, $\mathcal{N}M$ was a 2-D structure, i.e., it could always be parametrized by two parameters. In fact, $\mathcal{N}M$ is a two-dimensional continuous surface in a space of dimension 6.

For a smooth surface, we have introduced the function $\varphi_M(p,t) = p + tN_M(p)$, defined on $M \times (0, r_0)$. We now want to consider it as a function $\psi_M(p,n,t) = p + tn$ defined on $\mathcal{N}M \times (0, r_0)$. For smooth oriented surfaces, this is equivalent if the new definition of φ is restricted to the positive normal, i.e., $\mathcal{N}M^+$, since the latter is uniquely determined by m. But we have seen cases for which the normal was just partially specified by p, and this new definition of φ also applies to such situations. From now on, we assume that $M = \partial\Omega$ is the boundary of an bounded open subset of $\mathbb{R}^3$ and that (u, v) are local coordinates on $\mathcal{N}M^+$, so that we have a map $(u, v) \in U \mapsto (m(u, v), n(u, v))$ that locally parametrizes $\mathcal{N}M^+$ as a subset $\mathbb{R}^6$, U being an open subset of $\mathbb{R}^2$. We will let

$$A = \{(m(u, v), n(u, v)) : (u, v) \in U\}$$

denote the corresponding patch in $\mathcal{N}M^+$. We will assume that this map is differentiable in (u, v) (which may require the exclusion of some exceptional (negligible) sets from the integral that will be computed below. Our goal is to compute the volume of $V_r^+(M, A) = \psi_M(A \times (0, r))$.

The area form on $\mathcal{N}M^+$ is given by $g(u, v)\, du\, dv$ where

$$g(u, v)^2 = (|\partial_1 m|^2 + |\partial_1 n|^2)(|\partial_2 m|^2 + |\partial_2 n|^2) - (\partial_1 m^T \partial_2 m + \partial_1 n^T \partial_2 n)^2.$$

(In this expression, we have used $g(u, v) = \sqrt{EF - G^2}$, using (3.7), which still holds for two-dimensional surfaces in higher-dimensional spaces.)

One can check (we skip the proof) that the ratio $Q = |\det(\partial_1 m + t\partial_1 n_1, \partial_2 m + t\partial_2 n, n)|/g(u, v)$ is invariant under a change of local chart on $\mathcal{N}M^+$, allowing us to consider it as a function $Q(m, n)$ defined over A. When M is a smooth surface, this ratio can easily be computed, since we can assume that $\partial_1 m$ and $\partial_2 m$ correspond to the principal directions, yielding

$$Q = \frac{(1 + t\kappa_1)(1 + t\kappa_2)}{\sqrt{(1 + \kappa_1^2)(1 + \kappa_2^2)}}.$$

Returning to the general case, we have, by definition of Q:

$$\mathrm{Vol}(V_r^+(M, A)) = \int_U |\det(\partial_1 m + t\partial_1 n, \partial_2 m + t\partial_2 n, n)| du dv = \int_A Q d\sigma.$$

Assume that r can be chosen so that Q does not vanish for $t \in (0, r)$. (That such an r exists relates to the assumption of M having positive reach.) In this case, $\det(\partial_1 m + t\partial_1 n, \partial_2 m + t\partial_2 n, n)$ has constant sign, and one can expand Q in powers of t, yielding, for some functions S_0, S_1, S_2

$$\int_0^r \int_A Q d\sigma = r \int_A S_0 d\sigma - r^2 \int_A S_1 d\sigma + \frac{r^3}{3} \int_A S_2 d\sigma.$$

The functions S_k therefore provide densities for "generalized" curvature measures, defined on $\mathcal{N}M^+$ (instead of on M).

Recall that we assume that $M \subset \partial\Omega$ for some open set $\Omega \subset \mathbb{R}^2$. We can restrict the generalized curvature measures to M, letting for $B \subset \mathbb{R}^3$,

$$\mu_k(M, B) = \int_{\mathcal{N}M^+} \chi_B(m) S_k(m, n) d\sigma.$$

We now consider the case in which M is a triangulated surface and discuss the expression of this integral based on the location of the set B.

Face interiors: Let B be included in the interior of a face. Since M coincides there with a smooth surface with zero curvature, we have $\mu_0(M, B) = \mathrm{Area}(B)$, $\mu_1(M, B) = \mu_2(M, B) = 0$.

Convex edges: Now let B be included in the interior of a convex (salient) edge, e. At $p \in B$, normal vectors to Ω form the arc of the unit circle perpendicular to the edge delimited by the two outward normals to the neighboring faces. Fix an orientation of e and define $\beta_N(e)$ as the angle from the outward normal on the right to the one on the left of e (the saliency assumption implies that this angle is non-negative.) Now, on B, the normal bundle can be parametrized by $p = u(e/|e|) + n(v)$, where $n(v)$ is the normal to p that makes an angle v with one of the face normals. Using the fact that e and n are orthogonal, one finds $d(u, v) = 1$ and $|\det(\partial_u m, t\dot{n}_v, n)| = t$. This implies that $\mu_0 = \mu_2 = 0$ and $\mu_1 = \beta_N(e)\mathrm{length}(B)$.

Concave edges: The case of B included in a concave edge requires a little more care because Ω does not have positive reach on B. One can however split Ω into two parts on each side of the bisecting angle between the faces at e and apply the formula (letting Ω_1 and Ω_2 be the two sections)

$$\begin{aligned}\mu_k(\Omega, B) &= \mu_k(\Omega_1, B) + \mu_k(\Omega_2, B) - \mu_k(\Omega_1 \cap \Omega_2, B)\\ &= ((\pi - \beta_N)/2 + (\pi - \beta_N(e))/2 - \pi)\mathrm{length}(B)\\ &= -\beta_N \mathrm{length}(B),\end{aligned}$$

where $\beta_N(e)$ is again the absolute value of the angle between the normal to the faces meeting at e (taken between 0 and π).

Vertices: Let $B = \{v\}$, where v is a vertex of the triangulation. First note that, when M has positive reach, the volume of $V_r^+(M, \{v\})$ cannot be larger than that of the ball centered at v with radius r and is therefore an $O(r^3)$. From formula (4.8), this is also true when M is a triangulated surface. This implies that only the last term (the Gauss curvature measure) can be non-zero. The computation of this term is simplified if we also observe that the Gauss curvature measure does not change if we replace (locally at the vertex) Ω by $\Omega^c \cap \partial\Omega$, which corresponds to changing the orientation on M. This is because the function S_2 is an even function of the normal. Using this property, we get

$$2\mu_2(\Omega, B) = \mu_2(\Omega, B) + \mu_2(\Omega \cup \Omega^c, B) = \mu_2(\mathbb{R}^3, B) + \mu_2(M, B).$$

Since $\mu_2(\mathbb{R}^3, B) = 0$, it remains to compute $\mu_2(M, B)$. Let $F_1, \ldots, F_q$ represent the faces containing v. We want to apply the inclusion-exclusion formula to

$$\mu_2(M, B) = \mu_2\Big(\bigcup_{i=1}^{q} F_i, B\Big).$$

For this, we need to compute the Gauss curvatures in three special cases covering possible intersections between faces. The simplest case is $\mu_2(\{v\}, B)$. In this case, we can parametrize the normal bundle by $(m(u, u'), n(u, u')) = (v, n(u, u'))$, where (u, u') is a parametrization of the unit sphere, for which we assume that $\dot{n}_u$ and $\dot{n}_{u'}$ are orthogonal with unit norm. In this case, $d(u, u') = 1$ and $|\det(t\partial_1 n, t\partial_2 n, n)| = t^2$. This implies that $S_2 = 1$ and $\mu_2 = 4\pi$ (μ_2 is three times the volume of the sphere).

Now, let e be a segment having v as one of its extremities. Assume without loss of generality that e is supported by the first coordinate axis and $v = 0$. We can parametrize $\mathcal{N}e$ at v with $(u, u') \mapsto (v, n(u, u'))$, where $n(u, u')$ parametrizes the half-sphere that is normal to M at v. This provides $\mu_2(e, B) = 2\pi$.

The last case is a triangle F with vertex v. Let θ be the angle at v. Here, the normal bundle at v is the part of the unit sphere which is contained between the two hyperplanes normal to each edge of F incident at v, for which the volume is $2(\pi - \theta)/3$, so that $\mu_2(F, B) = 2(\pi - \theta)$.

Now, it remains to apply the inclusion-exclusion formula. This formula starts with $\sum_{i=1}^{q} \mu_2(F_i, B)$, which is $2q\pi - 2\sum_i \theta_i$. Then comes the sum of the measures associated with the intersection of two faces: this intersection is an edge for the q pairs of adjacent faces, and just $\{v\}$ for the $\binom{q}{2} - q$ remaining ones. This yields the contribution $2q\pi - 4\binom{q}{2}\pi$. We finally need to sum all terms for intersections of three or more sets, which is always equal to $\{v\}$. This is

$$4\pi \sum_{k\geq 3} (-1)^{k-1} \binom{q}{k} = 4\pi \left(1 - q + \binom{q}{2}\right),$$

where we used the fact that

$$\sum_{k\geq 0}(-1)^{k-1}\binom{q}{k} = (1-1)^q = 0.$$

Summing all the terms, we obtain $\mu_2(M, B) = 4\pi - 2\sum_{i=1}^{q}\theta_i$ so that

$$\mu_2(\Omega, B) = 2\pi - \sum_{i=1}^{q}\theta_i.$$

We have therefore obtained the curvature measures associated to an oriented triangulated surface [71]. For any set $B \in \mathbb{R}^3$, they are:

- The mean curvature measure:

  $$\mu_1(M, B) = \sum_{e\in E}\varepsilon_e\beta_N(e)\text{length}(B\cap e),$$

  where β_N is the angle between the normals to the faces at e in absolute value and $\varepsilon_e = 1$ if the edge is convex and -1 otherwise.
- The Gauss curvature measure:

  $$\mu_2(M, B) = \sum_{v\in V\cap B}\mu_2(M, v),$$

  with

  $$\mu_2(M, v_i) = 2\pi - \sum_{k\in\mathcal{F}_i}\theta_v(F_k).$$

Using these expressions, we can make approximations of the curvature at a given vertex by letting

$$\mu_1(M, B) \simeq |B|H_i \text{ and } \mu_2(M, B) \simeq |B|K_i$$

for a vertex v_i in the triangulation. Taking $B = A_i$, as defined in the previous section (see Eq. 4.1), we obtain the same approximation of the Gauss curvature as the one obtained from the discretization of the Gauss–Bonnet theorem. The formulas for the mean curvature differ, however.

4.2.4 *Discrete Gradient and Laplace–Beltrami Operators*

We conclude this section on triangulated surfaces with a computation of the discrete equivalent of the gradient and Laplacian on surfaces.

Let S be a triangulated surface, $\mathcal{V} = \{v_1, \ldots, v_M\}$ and $\mathcal{F} = \{F_1, \ldots, F_K\}$ denoting, respectively, the sets of vertices and faces of S. To simplify the discussion, we will assume that the surface has no boundary, i.e., each edge belongs to exactly two faces.

A function ψ defined on S assigns a value $\psi(v_i)$ to each vertex, and the gradient of ψ will be defined as a vector indexed over faces. To compute it, we first focus on a face, $F = F_k$ for some $k \in \{1, \ldots, K\}$ that we will drop from the notation until further notice. Let (v_1, v_2, v_3) be the vertices of F (ordered consistently with the orientation), and let $e_1 = v_3 - v_2$, $e_2 = v_1 - v_3$ and $e_3 = v_2 - v_1$. Let $c = (v_1 + v_2 + v_3)/3$ be the center of the face.

We define the gradient of ψ on F, denoted $\nabla_S \psi(F)$, as the gradient of the linear interpolation of ψ on F, i.e.,

$$\nabla_S \psi(F) = \nabla_F \hat{\psi}_F,$$

where $\hat{\psi}_F(a_1 v_1 + a_2 v_2 + a_3 v_3) = a_1 \psi(v_1) + a_2 \psi(v_2) + a_3 \psi(v_3)$ for $a_1 + a_2 + a_3 = 1$ and ∇_F is the gradient on the face F considered as a regular surface. From a computational viewpoint, $u = \nabla_S \psi(F)$ is such that $u = \alpha_1 e_1 + \alpha_2 e_2$ and $u^T(v_k - v_l) = \psi(v_k) - \psi(v_l)$ $(k, l = 1, 2, 3)$, which gives

$$\begin{aligned}\psi(v_3) - \psi(v_2) &= (\alpha_1 e_1 + \alpha_2 e_2)^T e_1,\\ \psi(v_1) - \psi(v_3) &= (\alpha_1 e_1 + \alpha_2 e_2)^T e_2.\end{aligned}$$

Let ψ_F be the column vector $[\psi(v_1), \psi(v_2), \psi(v_3)]^T$, M the 2 by 3 matrix

$$M = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \end{pmatrix}$$

and G_F the matrix

$$G_F = \begin{pmatrix} |e_1|^2 & e_1^T e_2 \\ e_1^T e_2 & |e_2|^2 \end{pmatrix}.$$

With this notation, the previous system is $M\psi_F = G_F \alpha$. We therefore have

$$u = [e_1, e_2]\alpha = [e_1, e_2] G_F^{-1} M \psi_F.$$

We first notice that $\det G_F = |e_1|^2 \, |e_2|^2 - (e_1^T e_2)^2 = (|e_1||e_2| \sin\theta_3)^2$, where θ_3 is the angle at v_3. It is therefore equal to $4a(F)^2$, where $a(F)$ is the area of F. Moreover, we can write:

$$\begin{aligned} G_F^{-1} M \psi_F &= \det(G_F)^{-1} \begin{pmatrix} |e_2|^2 & -e_1^T e_2 \\ -e_1^T e_2 & |e_1|^2 \end{pmatrix} \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \end{pmatrix} \psi_F \\ &= \det(G_F)^{-1} \begin{pmatrix} -e_2^T e_1 & -e_2^T e_2 & -e_2^T e_3 \\ e_1^T e_1 & e_1^T e_2 & e_1^T e_3 \end{pmatrix} \psi_F, \end{aligned} \quad (4.9)$$

in which we have used the identity $e_3 = -e_1 - e_2$. Introducing the vector

$$h_\psi(F) = \psi(v_1) e_1 + \psi(v_2) e_2 + \psi(v_3) e_3$$

and the matrix

$$D_F = e_2 e_1^T - e_1 e_2^T = e_1 e_3^T - e_3 e_1^T = e_3 e_2^T - e_2 e_3,$$

a little computation yields

$$\nabla_S \psi(F) = \frac{D_F}{4a(F)^2} h_\psi(F).$$

We now pass to the computation of the discrete Laplace–Beltrami operator, which we define via the discrete analog of the property

$$\int_S |\nabla_S \psi|^2 d\sigma_S = - \int_S \psi \Delta_S \psi d\sigma_S$$

that characterizes the operator on smooth surfaces without boundary. For triangulated surfaces, we will identify $\Delta_S \psi$ via

$$\sum_{k=1}^{K} |\nabla_S \psi(F_k)|^2 a(F_k) = - \sum_{i=1}^{N} \psi(v_i)(\Delta_S \psi)(v_i)|A_i|,$$

where $|A_i|$ is the area attributed to vertex v_i (using, for example, Voronoï cells).

For a given face F, we can write (using the previous notation): $|\nabla_S \psi(F)|^2 = \alpha^T G_F \alpha = \psi_F^T M^T G_F^{-1} M \psi_F$. Applying M^T to (4.9), we get

$$M^T G_F^{-1} M = \det(G_F)^{-1} \begin{pmatrix} |e_1|^2 & e_1^T e_2 & e_1^T e_3 \\ e_1^T e_2 & |e_2|^2 & e_2^T e_3 \\ e_1^T e_3 & e_2^T e_3 & |e_3|^2 \end{pmatrix}.$$

Let Σ_F denote this last matrix. We can write:

$$\sum_{k=1}^{K} \frac{\psi_{F_k}^T \Sigma_{F_k} \psi_{F_k}}{4a(F_k)} = \frac{1}{4} \sum_{i=1}^{N} \psi(v_i) \sum_{k \in \mathcal{F}_i} (|e_{ik}|^2 \psi(v_i) + e_{ik}^T e'_{ik} \psi(v'_{ik}) + e_{ik}^T e''_{ik} \psi(v''_i k))/a(F_k),$$

where v'_{ik} and v''_{ik} are the other two vertices of F_k, $k \in \mathcal{F}_i$ (in addition to v_i, ordered according to the orientation) and e_{ik}, e'_{ik} and e''_{ik} are, respectively, the edges opposed to v_i, v'_{ik} and v''_{ik} in F_k. This implies that one should define

$$\Delta_S \psi(v_i) = -\frac{1}{4|A_i|} \sum_{k \in \mathcal{F}_i} \left(|e_{ik}|^2 \psi(v_i) + e_{ik}^T e'_{ik} \psi(v'_{ik}) + e_{ik}^T e''_{ik} \psi(v''_{ik}) \right)/a(F_k).$$

One can rewrite this discrete Laplacian in terms of angles. Denoting as before by θ'_{ik} and θ''_{ik} the angles at v'_{ik} and v''_{ik}, one has

$$e_{ik}^T e'_{ik} = -\cos\theta''_{ik}|e_{ik}|\,|e'_{ik}| = -2\,\mathrm{ctn}\,\theta''_{ik}a(F_k).$$

Similarly, $e_{ik}^T e''_{ik} = -2\,\mathrm{ctn}\,\theta'_{ik}a(F_k)$ and, since the sum of the edges is 0,

$$|e_{ik}|^2 = -e_{ik}^T(e'_{ik} + e''_{ik}) = 2(\mathrm{ctn}\,\theta'_{ik} + \mathrm{ctn}\,\theta''_{ik})a(F_k).$$

One can therefore write

$$\Delta_S\psi(v_i) = \frac{1}{2|A_i|}\sum_{k\in\mathcal{F}_i}\big(\mathrm{ctn}\,\theta''_{ik}(\psi(v'_{ik}) - \psi(v_i)) + \mathrm{ctn}\,\theta'_{ik}(\psi(v''_{ik}) - \psi(v_i))\big),$$

which provides a discrete definition of the Laplace–Beltrami operator on S. This formula is sometimes called the "cotangent formula" [232].

4.3 Consistent Approximation

So far the concepts we have defined for triangulated surfaces have been directly inspired by the corresponding notions in the theory of smooth surfaces. Here, we provide some results evaluating how well a triangulated surface can approximate a smooth one, and whether quantities defined on triangulated surfaces are good estimates of the same quantities computed on the surface that is being approximated.

Because this analysis will be important for further purposes, we focus on the approximation of integrals $\int_\Sigma h(p)\,d\sigma_\Sigma(p)$ over a C^2 regular surface Σ by sums

$$\sum_{F\in\mathcal{F}} h(c_F)a(F),$$

where $\mathcal{F}$ is the set of faces of a triangulated surface S, c_F is the center of mass of face F and $a(F)$ its area.

To handle situations in which Σ is a surface with boundary, we also assume that another C^2 regular surface, Σ', is given, extending Σ so that $\Sigma \cup \partial\Sigma \subset \Sigma'$. If Σ is a closed surface, we can take $\Sigma' = \Sigma$. We let $\varphi : \Sigma' \times (-\rho, \rho) \to \mathbb{R}^3$ be the normal map, so that $\varphi(p,t) = p + tN_{\Sigma'}(p)$, and we assume that ρ is small enough that φ is a diffeomorphism onto its image, denoted by U. For $q \in U$, we let $\xi(q)$ be the closest point to q in Σ', i.e., the unique p such that $\varphi(p,t) = q$ for some $t \in \rho$. We note that

$$\xi(\varphi(p,t)) = p$$

for all p, t, so that $d\xi \circ \varphi\, \partial_p \varphi = \mathrm{Id}$, with $\partial_p \varphi = \mathrm{Id} + t dN_{\Sigma'}$. Fixing p and letting u_1, u_2 denote principal directions at p, with principal curvatures κ_1 and κ_2, we therefore have, for $q = \varphi(p, t)$

$$d\xi(q)u_i = \frac{u_i}{1 + t\kappa_i},\ i = 1, 2.$$

Similarly, $d\xi \circ \varphi\, \partial_t \varphi = 0$, so that $d\xi(q)N_{\Sigma'}(p) = 0$. We assume in the following that ρ is chosen small enough that $1 + t\kappa_i$ is bounded away from zero for $|t| \le \rho$ and $i = 1, 2$.

Let S be a triangulated surface, with the usual notation $\mathcal{V}, \mathcal{F}, \mathcal{E}$ for the sets of vertices, faces and edges in S. We will assume that $S \subset U$. For such a surface, we define the following constants:

- $\varepsilon_1(S, \Sigma') = \sup_{q \in S} |\xi(q) - p|$, the distance from S to Σ'.
- $\varepsilon_2(S, \Sigma') = 1 - \min_{F \in \mathcal{F}} \min_{q \in F} N_S(F)^T N_{\Sigma'}(\xi(q))$.
- $\delta(S) = \max_{(v_1, v_2) \in \mathcal{E}} |v_1 - v_2|$, the maximum edge size in S.
- $\varepsilon_3(S, \Sigma) = d_H(\xi(S), \Sigma)$, where d_H is the Hausdorff distance

$$d_H(A, A') = \max\left(\sup_{x \in A} \mathrm{dist}(x, A'), \sup_{x \in A'} \mathrm{dist}(x, A)\right).$$

We will also let $\varepsilon(S, \Sigma') = \max(\varepsilon_1, \varepsilon_2, \varepsilon_3, \delta)$, and we will say that a sequence of triangulations $S^{(n)}$ converge to Σ if $\varepsilon(S^{(n)}, \Sigma) \to 0$.

With this notation, we have the following theorem.

Theorem 4.3 *Let h be a continuous function on a compact neighborhood of U. Then*

$$\int_\Sigma h(p)\, d\sigma_\Sigma(p) = \sum_{f \in \mathcal{F}} h(c_f) a(f) + |S| O(\varepsilon). \tag{4.10}$$

Proof Consider $F \in \mathcal{F}$. Let v_1, v_2, v_3 denote the vertices of f and $c = (v_1 + v_2 + v_3)/3$. Let $e_{ij} = v_j - v_i$. We first compute $\int_{\xi(f)} h(p)\, d\sigma_{\Sigma'}(p)$, for which we can use the local chart $\psi : (x, y) \mapsto \xi(v_1 + xe_{12} + ye_{13})$, for $x, y > 0$, $x + y < 1$. Then, letting $\mathcal{T}$ denote the triangle $\{(x, y) : x, y > 0, x + y < 1\}$,

$$\int_{\xi(F)} h(p)\, d\sigma_{\Sigma'}(p) = \int_{\mathcal{T}} h(\psi(x, y)) |\partial_x \psi \times \partial_y \psi|\, dx\, dy.$$

Note that, using the same notation as above for principal directions and curvatures on Σ' at $p = \psi(x, y)$,

$$\partial_x \psi = d\xi \circ \psi\, e_{12} = \frac{e_{12}^T u_1}{1 + t\kappa_1} u_1 + \frac{e_{12}^T u_2}{1 + t\kappa_2} u_2$$
$$\partial_y \psi = d\xi \circ \psi\, e_{13} = \frac{e_{13}^T u_1}{1 + t\kappa_1} u_1 + \frac{e_{13}^T u_2}{1 + t\kappa_2} u_2$$

so that

$$\begin{aligned}\partial_x\psi \times \partial_y\psi &= \frac{e_{12}^T u_1\, e_{13}^T u_2 - e_{12}^T u_2\, e_{13}^T u_1}{(1+t\kappa_1(p))(1+t\kappa_2(p))} N_{\Sigma'}(p)\\ &= \frac{(e_{12}\times e_{13})^T N_{S'}(p)}{(1+t\kappa_1(p))(1+t\kappa_2(p))} N_{\Sigma'}(p)\\ &= 2a(F)\frac{N_F^T N_{S'}(p)}{1+2tH(p)+t^2K(p)} N_{\Sigma'}(p),\end{aligned}$$

where H and K denote the mean and Gauss curvatures. We therefore have

$$\int_{\xi(F)} h(p)\, d\sigma_{\Sigma'}(p) = 2a(F)\int_{\mathcal{T}} \frac{N_f^T N_{\Sigma'}(\psi(x,y))\, h(\psi(x,y))}{1+2tH(\psi(x,y))+t^2K(\psi(x,y))}\, dx\, dy,$$

from which one immediately deduces that

$$\begin{aligned}\left|\int_{\xi(F)} h(p)\, d\sigma_{\Sigma'}(p) - a(F)h(c)\right| &= \left|\int_{\xi(F)} h(p)\, d\sigma_{\Sigma'}(p) - 2a(F)\int_{\mathcal{T}} h(c)\, dx\, dy\right|\\ &\le a(F)\max_{p\in F}\left|h(c) - \frac{N_F^T N_{\Sigma'}(\xi(p))\, h(\xi(p))}{1+2tH(\xi(p))+t^2K(\xi(p))}\right|.\end{aligned}$$

Introduce the modulus of continuity of h

$$\omega_h(\eta) = \max_{x,y\in U, |x-y|<\eta} |h(x)-h(y)|.$$

One has,

$$\begin{aligned}&\left|h(c) - \frac{N_F^T N_{\Sigma'}(\xi(p))\, h(\xi(p))}{1+2tH(\xi(p))+t^2K(\xi(p))}\right| \le\\ &\qquad \omega_h(\delta+\varepsilon_1) + \|h\|_\infty\varepsilon_2 + \|h\|_\infty\varepsilon_1 \max_{|t|\le\varepsilon_1, p\in\Sigma'} \frac{2H(p)+tK(p)}{1+2tH(p)+t^2K(p)},\end{aligned}$$

so that

$$\left|\int_{\xi(F)} h(p)\, d\sigma_{\Sigma'}(p) - a(F)h(c)\right| = a(F)\, O(\varepsilon).$$

Summing over all faces, we obtain the fact that

$$\int_{\xi(S)} h(p)\, d\sigma_{\Sigma'}(p) = \sum_{F\in\mathcal{F}} h(c_F)a(F) + |S|O(\varepsilon).$$

It now suffices to write

$$\int_{\xi(S)} h(p)\, d\sigma_{\Sigma'}(p) = \int_{\Sigma} h(p)\, d\sigma_{\Sigma}(p) + O(\varepsilon_3)$$

to conclude the proof of (4.10). □

Note also that we can replace ε_2 by $\varepsilon_2' = 1 - \min_{F\in\mathcal{F}} N_S(F)^T N_{\Sigma'}(\xi(c_F))$ (or any other point in F), because $\|d\xi\|$ being bounded implies $|N_{\Sigma'}(\xi(q)) - N_{\Sigma'}(c_F)| = o(\delta)$ for $q \in F$.

Almost the same proof can be applied to sums involving normal vectors, yielding, for example

$$\int_{\Sigma} h(p)^T N_{\Sigma}(p)\, d\sigma_{\Sigma}(p) = \sum_{F\in\mathcal{F}} h(c_F)^T N_F a(F) + |S| o(\varepsilon) \tag{4.11}$$

for continuous vector-valued functions h.

Note also that, if $h : \mathbb{R}^3 \to \mathbb{R}$ is a C^1 function, one has

$$\max_F |\nabla_\Sigma h(\xi(c_F)) - \nabla_S h(F)| = o(\varepsilon).$$

Indeed, we have defined $\nabla_S h(F) = \nabla_F(\hat{h}_F)$, where $\hat{h}_F$ is a linear interpolation of h on F, but $\nabla_F(\hat{h}_F) = \nabla_F h(c_F) + o(\delta)$ because h is C^1. Moreover, letting ∇h denote the $\mathbb{R}^3$ gradient of h, we have

$$\nabla_F h(c_F) = \nabla h(c_F) - (N_F^T \nabla h(c_F)) N_F$$

and

$$\nabla_\Sigma h(\xi(c_F)) = \nabla h(\xi(c_F)) - (N_\Sigma(\xi(c_F))^T \nabla h(\xi(c_F))) N_\Sigma(\xi(c_F))$$

and these two quantities differ as $O(\varepsilon)$.

One can consider other approximations of geometric quantities and their convergence when triangulated surfaces approximate smooth ones with increasing accuracy. See, for example, [146, 211], in which an equivalence is shown between correct approximation of normals, or metric tensors, of area and of the Laplace–Beltrami operator.

4.4 Isocontours and Isosurfaces

To conclude this chapter, we discuss methods that compute shapes (curves or surfaces) from discrete image data. We will discuss approaches based on energy minimization in Chap. 5. Here, we focus on what is probably the simplest approach, which is to define curves of surfaces implicitly based on interpolation of the image values.

Assuming that the image f is defined on a discrete grid, we will interpolate it as a function $\hat{f} : D \subset \mathbb{R}^d \to \mathbb{R}$ and define a shape as the level set

$$S_\lambda = \{m : f(m) = \lambda\}$$

for a properly chosen threshold, λ. As we know, if the gradient of $\hat{f}$ does not vanish on S_λ, this provides a smooth curve or surface (or a union of such).

The concrete implementation of such an approach presents a few challenges, however. We will start with a discussion of the simpler two-dimensional case, which will help in addressing the computation of isosurfaces, which is more intricate.

4.4.1 Computing Isocontours

We consider here a two-dimensional grid, $\mathcal{G}$, which is formed by points $p(s,t) = (x_s, y_t)$, where $(x_s, s = 1, \ldots, M)$ is a discretization of the horizontal axis and $(y_t, t = 1, \ldots, N)$ a discretization of the vertical axis. We assume that a discretization of a smooth function f is observed, via the collection

$$(f_{st} = f(p(s,t)), s = 1, \ldots, M, t = 1, \ldots, N).$$

The problem is to compute the isocontour $(f = \lambda)$ for a given λ, in the form of a polygon or a union of polygons. Without loss of generality, we can and will assume $\lambda = 0$ in the following discussion.

Since the exact function f is not observed, some interpolation must be done, and we will use bilinear interpolation for this. This means that the true (unobserved) f will be replaced by the interpolation (that we still denote by f, with some abuse of notation) which is defined as follows. Let $C(s,t)$ denote the cell (square) with vertices $p(s+\varepsilon_1, t+\varepsilon_2)$, $\varepsilon_i \in \{0,1\}$, $i = 1, 2$. Then, for $p = x, y \in C(s,t)$, let

$$f(p) = \sum_{\varepsilon_1,\varepsilon_2=0}^{1} \prod_{i=1}^{2} \big(\varepsilon_i r_i(p) + (1-\varepsilon_i)(1-r_i(p))\big) f_{s+\varepsilon_1,t+\varepsilon_2}, \tag{4.12}$$

with $r_1(p) = x - x_s$, $r_2(p) = y - y_t$.

Obviously, the set $\{f = 0\}$ is the union of its intersections with each cell in the grid, so that we can restrict to these intersections. Within a cell, f is given by (4.12), and the set $\{f = 0\}$ is either empty, or a line segment, or one or two branches of a hyperbola. This is because, introducing the coordinates $\xi = (x - x_s)/(x_{s+1} - x_s)$ and $\eta = (y - y_s)/(y_{s+1} - y_s)$, we can rewrite $f(p)$ in the cell as (up to a positive multiplicative constant):

$$\begin{aligned} f(p) &= f_{--}(1-\xi)(1-\eta) + f_{+-}\xi(1-\eta) + f_{-+}(1-\xi)\eta + f_{++}\xi\eta \\ &= \rho\Big(\big(\xi + \frac{f_{-+} - f_{--}}{\rho}\big)\big(\eta + \frac{f_{+-} - f_{--}}{\rho}\big) - \frac{f_{++}f_{--} - f_{-+}f_{+-}}{\rho^2}\Big) \end{aligned}$$

if $\rho := f_{++} - f_{+-} - f_{-+} + f_{--} \neq 0$ and

$$f(p) = (f_{+-} - f_{--})\xi + (f_{-+} - f_{--})\eta + f_{--}$$

if $\rho = 0$. In this formula, f_{++}, f_{+-}, f_{-+} and f_{--} are the values of f at the vertices of the cell.

We will approximate the intersection by line segments intersecting the edges of the cell. There can be 0, 1 or 2 such line segments, and we now discuss when these situations occur. An important observation is that, because the bilinear interpolation is linear when restricted to the edges of the cell, there is at most one intersection of the set $\{f = 0\}$ with each edge, and this is only possible when f takes different signs at each of the edge end-points. When this occurs, the points on the edges at which $f = 0$ can be easily computed by solving a linear equation. They will form the vertices of the polygonal line. The following, the proof of which we skip, can be justified directly from the quadratic expression of f in the cell.

(a) If all f_{++}, f_{+-}, f_{-+} and f_{--} have the same sign: there is no intersection with the edges, and therefore no intersection with the cell.
(b) If three of the values have the same sign, the last one having the opposite sign, there are two vertices in the cell, and one edge connecting them.
(c) If two values have the same sign on one edge and two have the opposite sign on the opposite edge, here also, there are two vertices and one edge.
(d) If the function changes sign on all the edges, there are four vertices and two edges. There are two subcases, letting $\delta = f_{++}f_{--} - f_{-+}f_{+-}$.
    (i) If $\delta > 0$, then one edge links the vertex on $\{\xi = 0\}$ to the one on $\{\eta = 1\}$, and the other the vertex on $\{\eta = 0\}$ to the one on $\{\xi = 1\}$.
    (ii) If $\delta < 0$, then one edge links the vertex on $\{\xi = 0\}$ to the one on $\{\eta = 0\}$, and the other the vertex on $\{\eta = 1\}$ to the one on $\{\xi = 1\}$.

Cases (a), (b) and (c) can be decided based on the signs of f only. Case (d) is called ambiguous because it requires the exact numerical values of f. There are a few additional exceptional cases that are left aside in this discussion. When $f = 0$ at one of the vertices of the cell, this vertex is also in the polygonal line. It connects to other vertices at opposite edges of the cell, unless one of the cell edges that contain it is included in the polygon. There is no ambiguous situation in that case.

Case (d) with $\delta = 0$ is more of a problem, because it corresponds to a situation in which the interpolated surface is the intersection of two lines and therefore has a singular point. One cannot lift this ambiguity, and one of the options (i) and (ii) should be selected. The selection cannot be completely arbitrary because this could create holes in the reconstructed polygons. One possible rule is to take one option (say (i)) when $\rho > 0$ and the other one when $\rho < 0$. The combination of case (d) and

$\delta = \rho = 0$ implies that $f = 0$ at all vertices of the cell which therefore should be in the final polygon, but there is an unsolvable ambiguity as to how they should be connected.

There is another way to handle case (d), disregarding δ, based, as we just discussed, on the sign of ρ, yielding

(d)′ In case (d) above, take solution (i) if $\rho > 0$ and (ii) otherwise.

The resulting algorithm is simpler, because, in case (d), the sign of ρ can be computed directly based on the signs of f on the vertices of the cell. It does not correspond to the bilinear approximation anymore, but this approximation was somewhat arbitrary anyway. It does break the symmetry of the solution, in the sense that, if f is replaced by $-f$, the isocontours computed using (d)′ will differ. This is illustrated in Fig. 4.2

In addition to allowing for the segmentation of specific shapes from images, when the interior of the shape is, say, darker than its exterior, isocontours have been used as basic components of image processing algorithms that are contrast-invariant in the sense of mathematical morphology. A good introduction to this and to the related literature can be found in [53, 54].

Finally, let us note that isocontours can be easily oriented in accordance with our convention for implicit contours, by simply ensuring that grid points with negative values of f lie on the left of each oriented edge.

4.4.2 Computing Isosurfaces

We now pass to the case of level sets for functions defined over three dimensions, and describe the construction of triangulated isosurfaces. Although the problem is in principle similar to the two-dimensional case, the solution is notably more complex, mainly because of the large number of ambiguous situations in the determination of the boundary. There is indeed a large literature on the subject, and the reader can refer (for example) to [39] for a recent bibliography.

The three-dimensional generalization of the algorithm that we have presented for isocontouring is called *marching cubes* [178], and progressively builds a triangulation by exploring every grid cell on which the function changes sign. We will use a notation similar to the previous section, and let $\mathcal{G}$ be a regular three-dimensional grid, with grid coordinates $p(s, t, u) = (x_s, y_t, z_u)$ where $s = 1, \dots, M$, $t = 1, \dots, N$, $u = 1, \dots, P$. Denote by $f_{stu} = f(p(s, t, u))$ the observed values of f on the grid. Like in two dimensions, we assume that f extends to the continuum with a trilinear interpolation as follows: Let $C(s, t, u)$ denote the cube (cell) with vertices $p(s + \varepsilon_1, t + \varepsilon_2, q + \varepsilon_3)$, $\varepsilon_i \in \{0, 1\}$, $i = 1, 2, 3$. Then, for $p = x, y, z \in C(s, t, u)$, let

$$f(p) = \sum_{\varepsilon_1,\varepsilon_2,\varepsilon_3=0}^{1} \prod_{i=1}^{3} (\varepsilon_i r_i(p) + (1 - \varepsilon_i)(1 - r_i(p))) f_{s+\varepsilon_1, t+\varepsilon_2, q+\varepsilon_3}$$

with $r_1(p) = x - x_s$, $r_2(p) = y - y_t$, $r_3(p) = z - z_u$.

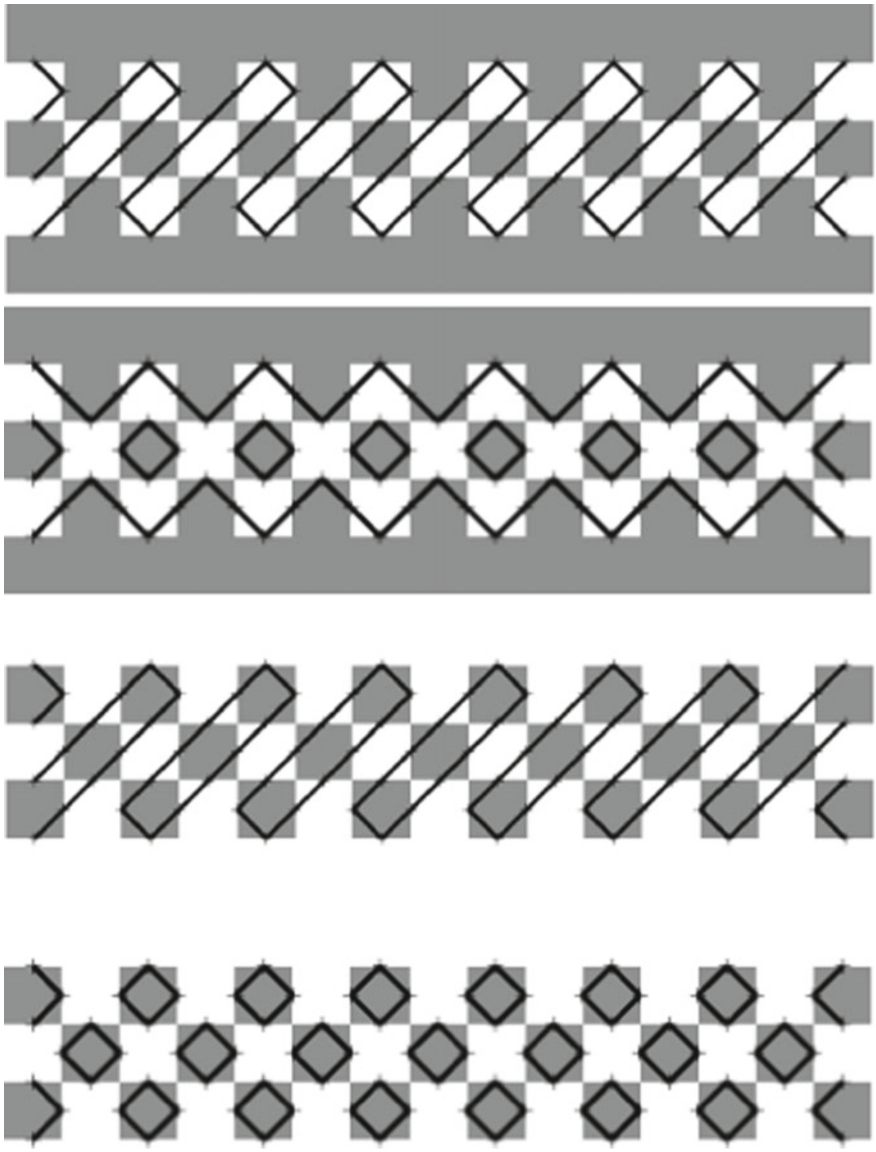

Fig. 4.2 Isocontouring a checkerboard strip using exact bilinear rule (d) (first row) and sign-based rule (d)' (second row). Note that the solutions are different, although both are plausible isocontours for the image. Gray levels are switched in the last two rows, without changing the solution for rule (d) (third row) and significantly altering it for rule (d)' (fourth row), yielding a third plausible solution

The determination of the vertices of the triangulation is similar to the two-dimensional case: the intersections of the level set $f = 0$ and the edges of the cubes $C(s, t, u)$ can be computed by solving a simple linear equation; on a given edge,

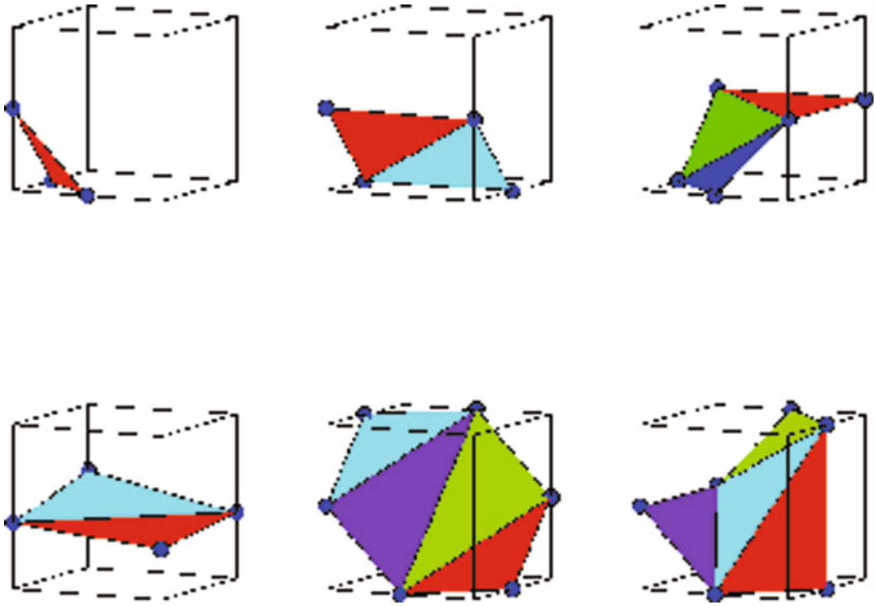

Fig. 4.3 Two-component (non-ambiguous) cases for the marching cubes algorithm

such an intersection exists only if f takes different signs at the end-points, and there can be at most one intersection. The difficulty is how to group these vertices into faces that provide a topologically consistent triangulation.

The main contribution of the marching cubes algorithm is to provide a method in which each cube is considered independently, yielding a reasonably simple implementation. The method works by inspection of the signs of f at the eight vertices of the cube. Like in two dimensions, there are some easy cases. The simplest is when all signs are the same, in which case the triangulation has no node on the cube. Other simple configurations are when the cube vertices of positive sign do not separate the other vertices in two or more regions and vice-versa. In this case, the triangulation has to separate the cube into two parts. There are, up to sign and space symmetry and up to rotation, six such cases, which are provided in Fig. 4.3.

Such triangulations can be efficiently described by labeling the vertices and the edges of the cube, as described in Fig. 4.4. We can describe a sign configuration on the cube by listing the vertices which have a positive sign. We can also describe each triangulation by listing, for each triangle, the three edges it intersects. Figure 4.3 therefore describes the six triangulations

$\{1\} : [(1, 4, 9)]$
$\{1, 2\} : [(2, 4, 9), (2, 4, 10)]$
$\{2, 5, 6\} : [(1, 2, 9), (2, 8, 9), (2, 8, 6)]$
$\{1, 2, 5, 6\} : [(2, 6, 4), (4, 6, 8)]$
$\{2, 3, 4, 7\} : [(1, 10, 6), (1, 6, 7), (1, 7, 4), (4, 7, 12)]$
$\{1, 5, 6, 7\} : [(1, 10, 11), (1, 11, 8), (8, 11, 7), (4, 1, 8)]$

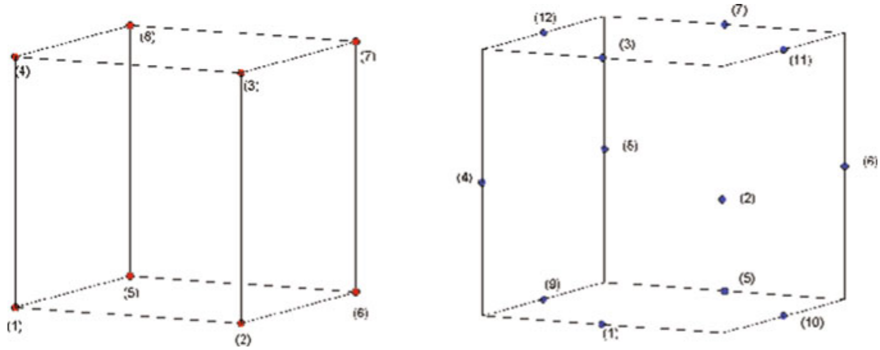

Fig. 4.4 Labels for the vertices (left) and edges (right) of the cube

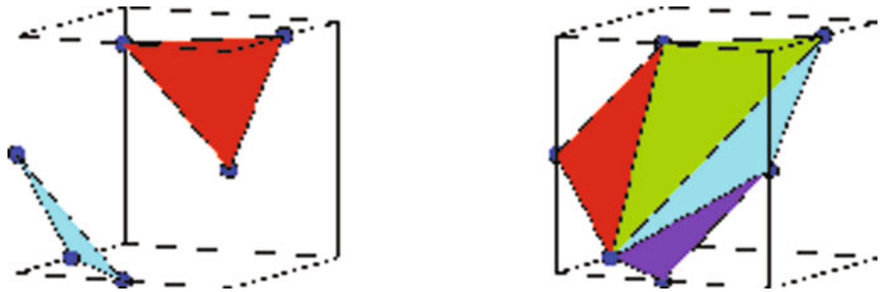

Fig. 4.5 Two triangulations associated to the {3, 8} sign configuration [(2, 3, 11), (1, 4, 9)] and [(4, 9, 3), (3, 9, 11), (9, 11, 2), (1, 2, 9)]

The cases when the signs form more than two connected components on the cube are problematic. They are ambiguous, because the way the surface crosses the cube cannot be decided from the sign pattern alone. One needs to rely on more information (i.e., the actual values of f at the nodes) to decide how to triangulate the surface within the cube, in order to avoid creating topological inconsistencies.

Take, for example, the case in which the cube vertices labeled (1) and (3) have signs distinct from the rest. Then, there are two possible ways (described in Fig. 4.5) in which the surface can cross the cube.

Another kind of ambiguous configuration is when two vertices in two opposite corners are isolated from the rest. Consider, for example, the situation when vertices 1 and 7 are positive while the rest are negative. Then the surface can do two things: either cut out the corners of the cube, or create a tunnel within the cube (see Fig. 4.6).

There have been successive attempts to improve the marching cubes algorithm from its original version ([178], in which the discussion was incomplete) [64, 209, 215, 218, 291] and untying the ambiguous cases. In addition to the two cases described in Figs. 4.5 and 4.6, five other ambiguous sign configurations can be listed, arising from combinations of these two basic cases. A complete description of all possible cases has been provided in [64], together with disambiguation rules. An

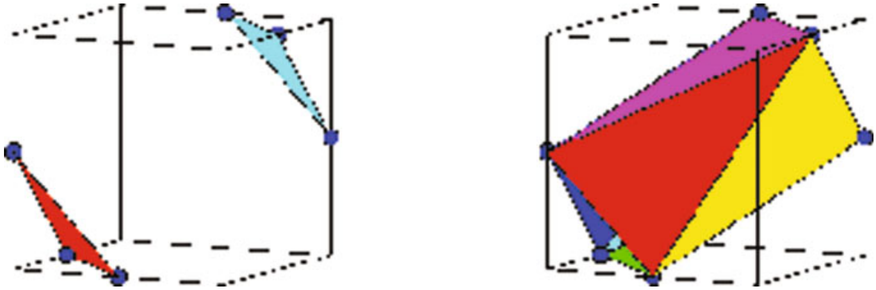

Fig. 4.6 Two triangulations associated to the {1, 7} sign configuration [(1, 4, 9), (7, 6, 11)] and [(1, 4, 11), (1, 11, 6), (1, 9, 6), (9, 6, 7), (9, 4, 7), (4, 11, 7)]

extensive theoretical and numerical analysis of the algorithm has been provided in [217] to which the reader is referred for complementary information, with the listing of all possible topologies within the cube.

If one drops the requirement to provide an accurate triangulation of zero-crossings of the linear interpolation of f within each cube, a reasonably simple option is available [209]. This approach has the disadvantage of breaking the sign-change invariance (which ensures that the computed triangulation should not change if f is

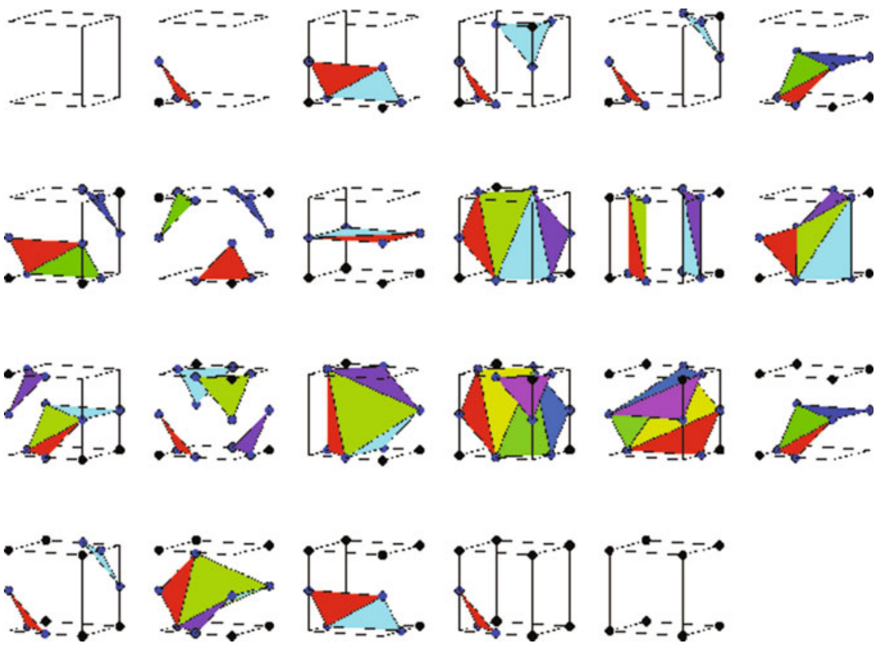

Fig. 4.7 Twenty-three configurations for consistent within-cube triangulation based on vertex signs. Dotted vertices correspond to positive values of the function

replaced by $-f$), but provides a very simple algorithm, still based on the signs of f on the vertices (it can be seen as a generalization of (d)$'$ in our discussion of the two-dimensional case). This results in 23 different cases (up to rotation invariance), listed in Fig. 4.7. This had to be compared to the 15 cases initially proposed in [178], which was invariant under sign change, but created topological errors.

An alternative to the marching cubes algorithm replaces cubic cells by tetrahedrons before computing the triangulation, which, when properly handled [57], provides a simpler and more stable procedure.

Extracting surfaces as level sets of functions is important even when the original data is not a three-dimensional image from which the region of interest is an isosurface. For example, when the original data is a set of unstructured points that roughly belong to the surface (i.e., they are subject to small errors) some of the commonly used algorithms that reconstruct the surface first reduce to the isosurface problem, trying to infer the signed distance function to the surface, at least in a neighborhood of the observed points. The approach used in [154] is to first approximate the tangent plane to the surface and then build the signed distance function. A similar goal is pursued in [11], using an approach based on computational topology.

Marching cubes (or tetrahedrons) have the drawback of providing a very large number of triangles, sometimes with very acute angles. Simplifying meshes is also the subject of a large literature, but this will not be addressed here (see, for example [100]).