Edge Contraction Algorithm

`Topics:` edge contraction, decimation, hierarchy, numerical error.

Edge contraction. The basic operation in simplifying a triangulated surface is the contraction of an edge. Let K be a pure 2-complex and assume for the moment that $|K|$ is a 2-manifold. The contraction of an edge $ab \in K$ removes ab together with the two triangles abx, aby and it mends the hole by gluing xa to xb and ya to yb as illustrated in Figure 1. Vertices a and b are

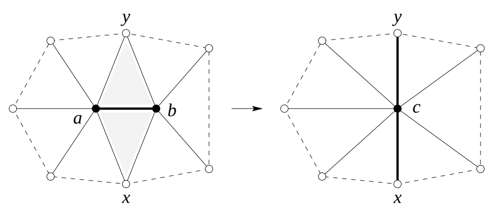

Figure 1: The contraction of edge ab. Vertices a and b are glued to a new vertex c.

glued to form a new vertex c. All simplices in the star of c are new, and the rest of the complex stays the same. To express this more formally we define the *cone* from a point x to a simplex τ as the union of line segments connecting x to points $p \in \tau$:

$$x \cdot \tau = \text{conv}(\tau \cup \{x\}).$$

It is defined only if x is not an affine combination of the vertices of τ. With this restriction, $x \cdot \tau$ is a simplex of one higher dimension: $\dim(x \cdot \tau) = 1 + \dim \tau$. For a *set* of simplices the cone is defined if it is defined for each simplex, and in this case $x \cdot T = \{x \cdot \tau \mid \tau \in T\}$. We also need generalizations of the star and the link from a single simplex to a set of simplices. Denote the closure without the (-1)-simplex as $\overline{T} = \text{Cl}\, T - \{\emptyset\}$. The *star* and *link* of T are

$$\begin{aligned} \text{St}\, T &= \{\sigma \in K \mid \sigma \geq \tau \in T\}, \\ \text{Lk}\, T &= \text{Cl}\,\text{St}\, T - \text{St}\, \overline{T}. \end{aligned}$$

For closed sets T the link is simply the boundary of the closed star. For example in Figure 1 the link of the set $\overline{ab} = \{ab, a, b\}$ is the cycle of dashed edges and hollow vertices bounding the closed star of $\overline{ab}$. The *contraction* of the edge ab is the operation that changes K to

$$L = K - \text{St}\, \overline{ab} \cup c \cdot \text{Lk}\, \overline{ab}.$$

This definition applies generally and does not assume that K is a manifold.

Decimation. The surface represented by K is simplified by performing a sequence of edge contractions. To get a meaningful result we prioritize the contractions by the numerical error they introduce. Contractions that change the topological type of the surface are rejected. Initially, all edges are evaluated and stored in a priority queue. The process continues until the number of vertices shrinks to the target number m. Let $n \geq m$ be the number of vertices in K.

```
while n > m and priority queue non-empty do
  extract top edge ab from priority queue;
  if contracting ab preserves topology then
    contract ab; n--
  endif
endwhile.
```

The priority queue takes time $O(\log n)$ per operation. Besides extracting the edge whose contraction causes the minimum error we remove edges that no longer belong to the surface and we add new edges. The number of edges removed and added during a single contraction is usually bounded by a small constant, but in the worst case it can be as large as $n-1$. Before performing an edge contraction we test whether or not it preserves the topological type of the surface. This is done by checking all edges and vertices in the link of $\overline{ab}$. Precise conditions to recognize edge contractions that preserve the type will be discussed in the next lecture.

Hierarchy. We visualize the actions of the algorithm by drawing the vertices as the nodes of an upside-down forest. The contraction of the edge ab combines vertices a and b into a new vertex c. In the forest this is reflected by introducing c as a new node and declaring it the parent of a and b. The leaves of the forest are the vertices of K, and the roots are the vertices of the decimated complex L, see Figure 2. We define a function

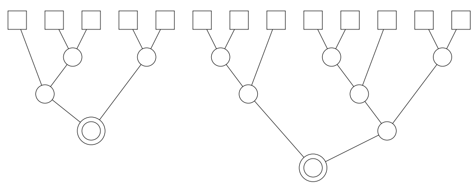

Figure 2: Vertices of K are shown as square nodes, intermediate vertices as circle nodes, and vertices of the final complex L as double circle nodes.

$g : \mathrm{Vert}\, K \to \mathrm{Vert}\, L$ that maps each vertex $u \in K$ to the root $g(u)$ of the tree in which u is a leaf. The preimage of a vertex $v \in L$ is the set of leaves $g^{-1}(v) \subseteq K$ of the tree with root v. The preimages of the roots partition the set of leaves:

$$\mathrm{Vert}\, K = \bigcup_{v \in L} g^{-1}(v),$$

where the union is over a collection of pairwise disjoint sets. Later, we will extend function g from vertices to edges and triangles. This will be useful in the study of structural connections between the surfaces K and L.

Numerical error. As mentioned above, a vertex $v \in \mathrm{Vert}\, L$ represents a subset $g^{-1}(v) \subseteq \mathrm{Vert}\, K$ of the vertices in K. It makes sense to measure the numerical error at v by comparing v to the part of the original surface it represents. Specifically, we define the error at v as the sum of square distances of v from the planes spanned by triangles in the star of $g^{-1}(v)$. See Figure 3, which shows a vertex $v \in L$ and the triangles in the star of $g^{-1}(v)$. The preimage of v is the collection of seven solid vertices in the right half of the figure. The star of the preimage contains the five shaded triangles and the ring of white triangles around them. The shaded triangles have all their vertices in $g^{-1}(v)$ and the white triangles have either one or two vertices in the preimage.

Let H_v be the set of planes spanned by triangles in $\mathrm{St}\, g^{-1}(v)$. The sum of square distances is defined for every point $x \in \mathbb{R}^3$, so we can think of the error measure as a function $E_v : \mathbb{R}^3 \to \mathbb{R}$. Specific properties of this function will be discussed in the third lecture after this one. For now we just observe that the error function

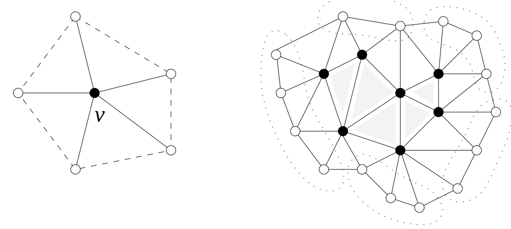

Figure 3: Vertex v and its star to the left and the corresponding piece of K to the right. The solid vertices on the right are preimages of v and the hollow vertices are preimages of the neighbors of v.

for the union of two sets of planes can be computed by inclusion-exclusion. Specifically, if $H_c = H_a \cup H_b$ then

$$E_c(x) = E_a(x) + E_b(x) - E_{ab}(x),$$

where E_{ab} is the error function defined by the intersection, $H_{ab} = H_a \cap H_b$. This formula together with a compact representation of error functions will be the mechanism we use to compute the error at a new vertex c. Given a set of planes there is generally a unique point that minimizes the corresponding error function. Instead of computing c directly from a and b we first construct $H_c = H_a \cup H_b$ and second choose c at the minimum of the error function E_c defined by H_c.

Bibliographic notes. The idea of using edge contractions for surface simplification appears first in Hoppe et al. [3]. They select contractions together with other local surface modification operations in an attempt to optimize a measure of distance between the original and the decimated surface. Hoppe [2] revisits the idea and shows how to use a given sequence of contractions for efficiently switching back and forth between representations on different levels of detail. The algorithm in these notes selects contractions greedily using the quadratic error measure as suggested by Garland and Heckbert [1].

[1] M. Garland and P. S. Heckbert. Surface simplification using quadratic error metrics. *Comput. Graphics*, Proc. SIGGRAPH 1997, 209–216.

[2] H. Hoppe. Progressive meshes. *Comput. Graphics*, Proc. SIGGRAPH 1996, 99–108.

[3] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald and W. Stützle. Mesh optimization. *Comput. Graphics*, Proc. SIGGRAPH 1993, 19–26.

Preserving Topology

Topics: manifolds, manifolds with boundary, open books, boundary, 2-complexes.

Manifolds. Suppose K is a 2-complex that triangulates a 2-manifold. Then every point $x \in |K|$ has a neighborhood homeomorphic to an open disk. To avoid lengthy sentences we just say the neighborhood *is* an open disk. This implies that in particular the star of every vertex u is an open disk. Strictly speaking this statement makes sense only if we replace the star by its underlying space, which we define as the union of simplex interiors, which is the set difference between the underlying spaces of two complexes:

$$\begin{aligned} |\mathrm{St}\, u| &= \bigcup_{\tau \in \mathrm{St}\, u} \mathrm{int}\, \tau \\ &= |\mathrm{Cl}\,\mathrm{St}\, u| - |\mathrm{Cl}\,\mathrm{St}\, u - \mathrm{St}\, u|. \end{aligned}$$

As it turns out the condition on vertex stars is sufficient to guaranteed that $|K|$ is a 2-manifold.

CLAIM 1. $|K|$ is a 2-manifold iff $|\mathrm{St}\, u| \approx \mathbb{R}^2$ for every vertex $u \in K$.

Now consider the contraction of an edge ab of K. Whether or not the contraction preserves the topological type depends on how the links of a and b meet. On a 2-manifold the link of each vertex is a circle. In Figure 4 to the left the two circles intersect in two points and the contraction preserves the topological type. To the right the circles intersect in a point and an edge, and in this case the contraction pinches the manifold along a newly formed edge which forms the base of a fin similar to the one in Figure 7. The condition that distinguishes topology preserving edge contractions from others is that the vertex links intersect in the link of the edge.

THEOREM 2A. Let K be the triangulation of a 2-manifold. The contraction of $ab \in K$ preserves the topological type iff $\mathrm{Lk}\, a \cap \mathrm{Lk}\, b = \mathrm{Lk}\, ab$.

A proof of the sufficiency of the link condition will be given in the next lecture.

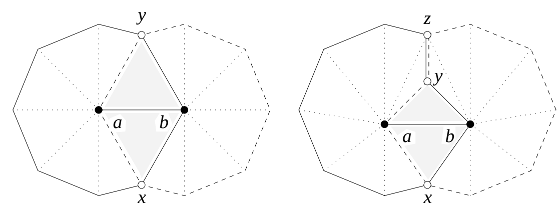

Figure 4: The edges of the link of a are solid and those of the link of b are dashed.

Manifolds with boundary. A triangulation K of a manifold with non-empty boundary also has vertices whose stars are open half-disks: $|\mathrm{St}\, u| \approx \mathbb{H}^2$. To keep the number of cases small we add a dummy vertex, ω, and the cone from ω to each boundary circle. This idea is illustrated in Figure 5. The boundary of $|K|$ consists

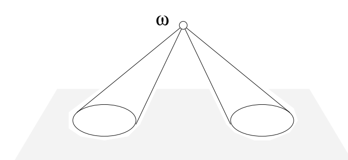

Figure 5: The two holes in the manifold are filled by adding the cone from ω to the circles bounding the holes.

of $\ell \geq 1$ circles triangulated by cycles $C_i \subseteq K$. We fill the holes by adding the cone from ω to every cycle:

$$K^\omega = K \cup (\omega \cdot \bigcup_{i=1}^{\ell} C_i).$$

In K^ω every vertex star is an open disk except possibly the star of ω. We denote the link of a vertex u in K^ω as $\mathrm{Lk}^\omega u$. The condition that distinguishes topology preserving edge contractions from others is now the same as for manifolds.

THEOREM 2B. Let K be the triangulation of a 2-manifold with boundary. The contraction of $ab \in$

K preserves the topological type iff $\mathrm{Lk}^\omega a \cap \mathrm{Lk}^\omega b = \mathrm{Lk}^\omega ab$.

The proof of this result is only mildly more complicated than that of the weaker Theorem 2A.

Open books. To attack the problem for general 2-complexes we need a better understanding of the different types of neighborhoods that are possible. We classify stars using a new type of space. The *open book with p pages* is the topological space $\mathbb{K}^2_p$ homeomorphic to the union of p copies of $\mathbb{H}^2$ glued along the common boundary line. For example, the open book with one page is the open half-disk and the open book with two pages is the open disk. The *order* of a simplex $\tau \in K$ is

$$\mathrm{ord}\,\tau = \begin{cases} 0 & \text{if } |\mathrm{St}\,\tau| \approx \mathbb{R}^2, \\ 1 & \text{if } |\mathrm{St}\,\tau| \approx \mathbb{K}^2_p,\ p \neq 2, \\ 2 & \text{otherwise.} \end{cases}$$

Figure 6 illustrates the definition with sketches of four vertex stars. The order of an edge in a 2-complex can only be 0 or 1, and the order of a triangle is always 0.

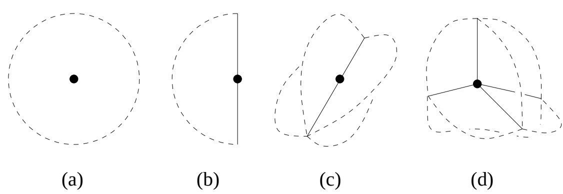

Figure 6: The underlying space of the vertex star in (a) is an open disk, in (b) is an open half-disk, in (c) is an open book with 4 pages, and in (d) is not an open book. The corresponding order of the vertex is 0 in (a), 1 in (b), 1 in (c), and 2 in (d).

Boundary. We generalize the notion of boundary in such a way that only triangulations of 2-manifolds have no boundary. At the same time we use the order information to distinguish between different types of boundaries. Specifically, the *j-th boundary* of a 2-complex K is the collection of all simplices with order j or higher:

$$\mathrm{Bd}_j K = \{\sigma \in K \mid \mathrm{ord}\,\sigma \geq j\}.$$

As an example consider the shark fin complex shown in Figure 7. It is constructed by gluing two closed disks along a simple path. This path is a contiguous piece of the boundary of one disk (the fin) and it lies in the interior of the other disk. Note that $|K|$ is a 2-manifold

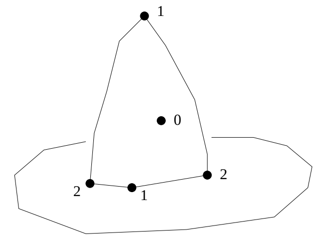

Figure 7: The shark fin 2-complex. A few of the vertices are high-lighted and marked with their order.

with boundary iff $\mathrm{Bd}_1 K = \mathrm{Bd}_2 K = \emptyset$. The 2-nd boundary of a 2-manifold with boundary is empty, but there are other spaces with this property. For example, the sphere together with its equator disk has empty 2-nd boundary. Its 1-st boundary is a circle of edges and vertices (the equator) whose stars are open books of 3 pages each.

2-complexes. We are now ready to study conditions under which an edge contraction in a general 2-complex preserves the topological type of that complex. As it turns out there does not exist a local condition that is sufficient and necessary, but there is a characterizing local condition for a more restrictive notion of type preservation. Let L be the 2-complex obtained from K by contracting an edge $ab \in K$. A *local unfolding* is a homeomorphism $f : |K| \to |L|$ that differs from the identity only outside the star of $\overline{ab}$, that is, $f(x) = x$ for all $x \in |K - \mathrm{St}\,\overline{ab}|$. The condition refers to links in $K^\omega = K \cup (\omega \cdot \mathrm{Bd}_1 K)$ and in $G^\omega = \mathrm{Bd}_1 K \cup (\omega \cdot \mathrm{Bd}_2 K)$. We denote the link of a simplex τ in K^ω by $\mathrm{Lk}^\omega_0 \tau$ and the link of τ in G^ω by $\mathrm{Lk}^\omega_1 \tau$.

THEOREM 2C. Let K be a 2-complex, ab an edge of K, and L the complex obtained by contracting ab. There is a local unfolding $|K| \to |L|$ iff

(i) $\mathrm{Lk}^\omega_0 a \cap \mathrm{Lk}^\omega_0 b = \mathrm{Lk}^\omega_0 ab$ and

(ii) $\mathrm{Lk}^\omega_1 a \cap \mathrm{Lk}^\omega_1 b = \emptyset$.

Instead of proving Theorem 2C, which is a bit tedious in any case, we show that there cannot be a similar condition that recognizes the existence of a general homeomorphism $|K| \to |L|$. The example we use is the folding chair complex displayed in Figure 8. Before the contraction of ab it consists of five triangles in the star of x and four disks U, V, Y, Z glued to the link of x. Vertices a and b belong to the 1-st boundary, but ab does not. It follows that ω violates condition (i) of Theorem 2C and there is therefore no local unfolding from $|K|$

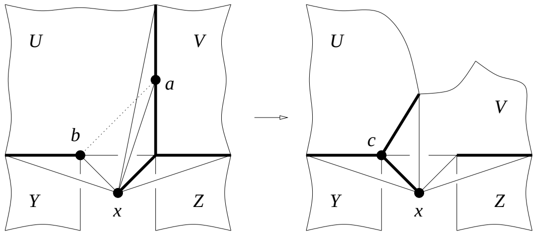

Figure 8: The folding chair complex. The bold edges belong to three triangles each.

to $|L|$. After the contraction there is one less triangle in the star of x, U loses two triangles, and V, Y, Z are unchanged. The contraction of ab exchanges left and right in the asymmetry of the complex. We can find a homeomorphism $|K| \to |L|$ that acts like a mirror and maps U to V, V to U, Y to Z, Z to Y. The homeomorphism is necessarily global and to detect it we can force any algorithm to look at every triangle of K.

Bibliographic notes. The material of this lecture is taken from a recent paper by Dey et al. [1]. It studies edge contraction in general simplicial complexes and proves results for 2- and for 3-complexes. The order of a simplex has already been defined in 1960 by Whittlesey [4], although in different words and notation. He uses the concept to study the topological classification of 2-complexes. O'Dunlaing et al. [2] use his results to show that deciding whether or not two 2-complexes have the same topological type is just as hard as deciding whether or not two graphs are isomorphic. No polynomial time algorithm is known, but it is also not known whether the graph isomorphism problem is NP-complete [3].

[1] T. K. Dey, H. Edelsbrunner, S. Guha and D. V. Nekhayev. Topology preserving edge contraction. Report rgi-tech-98-018, Raindrop Geomagic, Research Triangle Park, North Carolina, 1998.

[2] C. O'Dunlaing, C. Watt and D. Wilkins. Homeomorphism of 2-complexes is equivalent to graph isomorphism. Rept. TCDMATH 98-04, Math. Dept., Trinity College, Ireland, 1998.

[3] M. R. Garey and D. S. Johnson. *Computers and Intractibility. A Guide to the Theory of NP-Completeness.* Freeman, San Francisco, California, 1979.

[4] E. F. Whittlesey. Finite surfaces: a study of finite 2-complexes. *Math. Mag.* **34** (1960), 11–22 and 67–80.

Error Measure

Topics: signed distance, fundamental quadric, error, eigenvalues and eigenvectors.

Signed distance. The surface simplification algorithm measures the error of an edge contraction as the sum of square distances of a point from a collection of planes. A plane with unit normal vector v_i and offset δ_i contains all points p with orthogonal projection $\delta_i \cdot v_i$:

$$h_i = \{p \in \mathbb{R}^3 \mid p^T \cdot v_i = -\delta_i\},$$

see Figure 1. The signed distance of a point $x \in \mathbb{R}^3$ from the plane h_i is

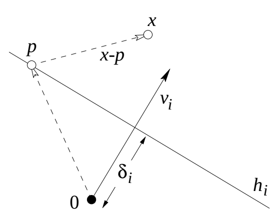

Figure 1: We use the unit normal vector to define the signed distance from h_i so v_i points from the negative to the positive side.

$$\begin{aligned} d(x, h_i) &= (x - p)^T \cdot v_i \\ &= x^T \cdot v_i + \delta_i \\ &= \mathbf{x}^T \cdot \mathbf{v}_i, \end{aligned}$$

where $\mathbf{x}^T = (x^T, 1)$ and $\mathbf{v}_i{}^T = (v_i{}^T, \delta_i)$. In words, the signed distance in $\mathbb{R}^3$ can be expressed as a scalar product in $\mathbb{R}^4$ as illustrated in Figure 2.

Fundamental quadric. The sum of square distances of a point x from a collection of planes H is

$$\begin{aligned} E_H(x) &= \sum_{h_i \in H} d^2(x, h_i) \\ &= \sum_{h_i \in H} (\mathbf{x}^T \cdot \mathbf{v}_i) \cdot (\mathbf{v}_i{}^T \cdot \mathbf{x}) \\ &= \mathbf{x}^T \cdot (\sum_{h_i \in H} \mathbf{v}_i \cdot \mathbf{v}_i{}^T) \cdot \mathbf{x}, \end{aligned}$$

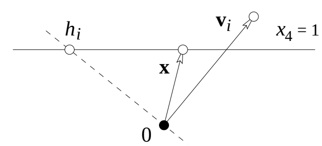

Figure 2: The 3-dimensional space $x_4 = 1$ is represented by the horizontal line. It contains point x and plane h_i, which in the 1-dimensional representation are both points.

where

$$\mathbf{Q} = \sum \mathbf{v}_i \cdot \mathbf{v}_i{}^T = \begin{pmatrix} A & B & C & D \\ B & E & F & G \\ C & F & H & I \\ D & G & I & J \end{pmatrix}.$$

is a symmetric 4-by-4 matrix referred to as the *fundamental quadric* of the map $E_H : \mathbb{R}^3 \to \mathbb{R}$. The sum of square distances is non-negative, so $\mathbf{Q}$ is positive semi-definite. The error of an edge contraction is obtained from an error function like $E = E_H$. Let $\mathbf{x}^T = (x_1, x_2, x_3, 1)$ and note that

$$\begin{aligned} E(x) &= \mathbf{x}^T \cdot \mathbf{Q} \cdot \mathbf{x} \\ &= Ax_1^2 + Ex_2^2 + Hx_3^2 \\ &\quad + 2(Bx_1x_2 + Cx_1x_3 + Fx_2x_3) \\ &\quad + 2(Dx_1 + Gx_2 + Ix_3) \\ &\quad + J. \end{aligned}$$

We see that E is a quadratic map that is non-negative and unbounded. Its graph can only be an elliptic paraboloid as illustrated in Figure 3. In other words, the preimage of a constant error value ϵ, $E^{-1}(\epsilon)$, is an ellipsoid. Degenerate ellipsoids are possible, such as cylinders with elliptic cross-sections and pairs of planes.

Error. The *error* of the edge contraction $ab \to c$ is the minimum value of $E(x) = E_H(x)$ over all $x \in$

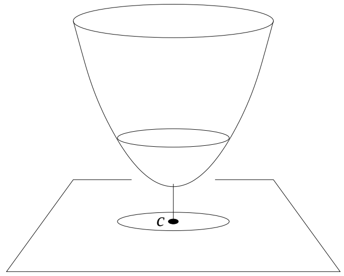

Figure 3: Illustration of $E = E_H$ in one lower dimension. The cross-section at a fixed height ϵ is an ellipse.

$\mathbb{R}^3$, where H is the set of planes spanned by triangles in the preimage of the star of the new vertex c. The geometric location of c is the point x that minimizes E. In the non-degenerate case this point is unique and can be computed by setting the gradient $\nabla E = (\partial E/\partial x_1, \partial E/\partial x_2, \partial E/\partial x_3)$ to zero. The derivative with respect to x_i is

$$\begin{aligned} \frac{\partial E}{\partial x_i}(x) &= \frac{\partial \mathbf{x}^T}{\partial x_i} \cdot \mathbf{Q} \cdot \mathbf{x} + \mathbf{x}^T \cdot \mathbf{Q} \cdot \frac{\partial \mathbf{x}}{\partial x_i} \\ &= \mathbf{Q}_i^T \cdot \mathbf{x} + \mathbf{x}^T \cdot \mathbf{Q}_i, \end{aligned}$$

where $\mathbf{Q}_i^T$ is the i-th row of $\mathbf{Q}$. The point $c \in \mathbb{R}^3$ that minimizes $E(x)$ is the solution to the system of three linear equations $Q \cdot x = q$, where

$$Q = \begin{pmatrix} A & B & C \\ B & E & F \\ C & F & H \end{pmatrix} \text{ and } q = \begin{pmatrix} D \\ G \\ I \end{pmatrix}.$$

Hence $c = Q^{-1} \cdot q$, and the sum of square distances of c from the planes in H is $E(c)$. The equation for c sheds light on the possible degeneracies. The non-degenerate case corresponds to $\operatorname{rank} Q = 3$, the case of a cylinder corresponds to $\operatorname{rank} Q = 2$, and the case of two parallel planes corresponds to $\operatorname{rank} Q = 1$. Rank 0 is not possible because Q is the non-empty sum of products of unit vectors.

Eigenvalues and eigenvectors. We may translate the planes by $-c$ so E attains its minimum at the origin. In this case $D = G = I = 0$ and $J = E(0)$. The shape of the ellipsoid $E^{-1}(\epsilon)$ can be described by the eigenvalues and eigenvectors of Q. By definition, the *eigenvectors* are non-zero vectors x that satisfy $Q \cdot x = \lambda \cdot x$. The value of λ is the corresponding *eigenvalue*. The eigenvalues are the roots of the *characteristic polynomial* of Q, which is

$$\begin{aligned} P(\lambda) &= \det \begin{pmatrix} A-\lambda & B & C \\ B & E-\lambda & F \\ C & F & H-\lambda \end{pmatrix} \\ &= \det Q - \lambda \cdot \operatorname{dtr} Q + \lambda^2 \cdot \operatorname{tr} Q - \lambda^3, \end{aligned}$$

where $\det Q$ is the determinant, $\operatorname{dtr} Q$ is the sum of cofactors of the three diagonal elements, and $\operatorname{tr} Q$ is the trace of Q. For symmetric positive semi-definite matrices the characteristic polynomial has three non-negative roots, $\lambda_1 \geq \lambda_2 \geq \lambda_3 \geq 0$. Once we have an eigenvalue we can compute the corresponding eigenvector to span the nullspace of the underconstrained system $(Q - \lambda) \cdot x = 0$.

What is the geometric meaning of eigenvectors and eigenvalues? For symmetric matrices the eigenvectors are pairwise orthogonal and can be viewed as defining another coordinate system for $\mathbb{R}^3$. The three symmetry planes of the ellipsoid $E^{-1}(\epsilon)$ coincide with the coordinate planes of this new system, see Figure 4. We can write the error function as

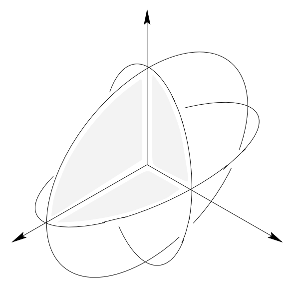

Figure 4: The ellipsoid is indicated by drawing the elliptic cross-sections along the three symmetry planes spanned by the eigenvectors.

$$\begin{aligned} E(x) &= \mathbf{x}^T \cdot \begin{pmatrix} \lambda_1 & & & \\ & \lambda_2 & & \\ & & \lambda_3 & \\ & & & J \end{pmatrix} \cdot \mathbf{x} \\ &= \lambda_1 x_1^2 + \lambda_2 x_2^2 + \lambda_3 x_3^3 + J. \end{aligned}$$

The preimage for a fixed error $\epsilon > J$ is the ellipsoid with axes of half-lengths $\sqrt{(\epsilon - J)/\lambda_i}$ for $i = 1, 2, 3$.

Bibliographic notes. The idea of using the sum of square distances from face planes for surface simplification is due to Garland and Heckbert [1]. Eigenvalues and eigenvectors of matrices are topics in linear algebra. A very readable introductory text is the book by Gilbert Strang [2].

[1] M. Garland and P. S. Heckbert. Surface simplification using quadratic error metrics. *Computer Graphics*, Proc. siggraph 1997, 209–216.

[2] G. Strang. *Introduction to Linear Algebra.* Wellesley-Cambridge Press, Wellesley, Massachusetts, 1993.