Chapter 6

Conformal Maps

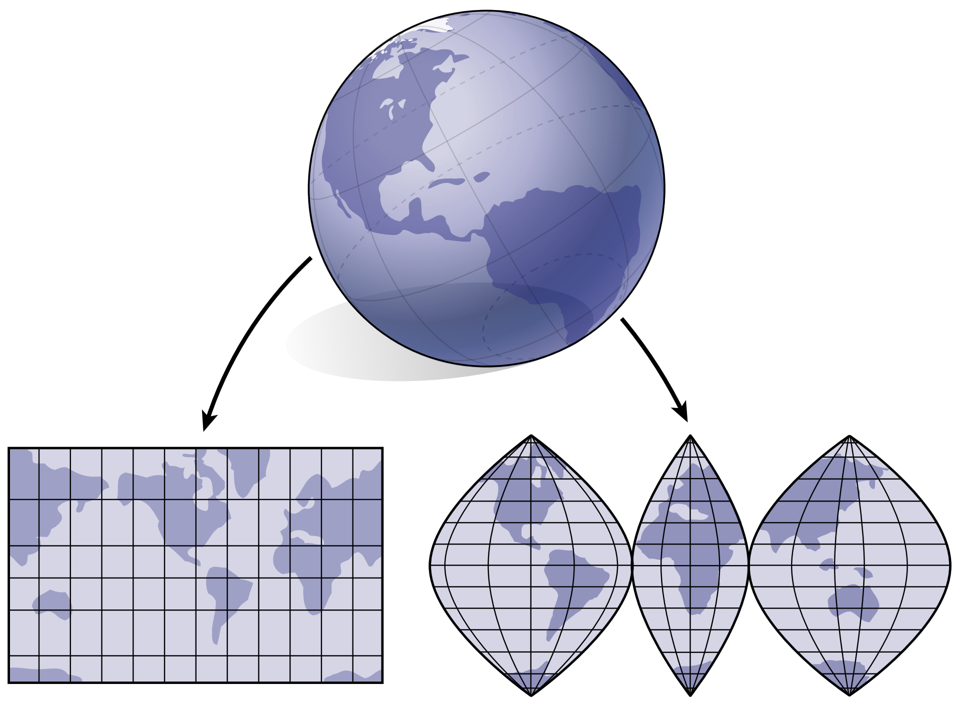

Figure 6.1: *Finding nice charts (or conversely, parameterizations) for curved surfaces is the cartography problem.*

Charts and their inverse, parameterizations, are of fundamental importance when dealing with surfaces. They allow us to pull back quantities from a canonical domain, a region in $\mathbb{R}^2$, or describe the surface as a function over some region in $\mathbb{R}^2$ to begin with. Coming up with such maps from a curved surface to a flat domain is of course the ancient cartography problem (Fig. 6.1). If the original surface has non-zero Gaussian curvature, say, for example, the earth, such a map cannot be free of distortion. But there is leeway as to what type of distortion we want to admit. A special class of such controlled distortion maps are the **conformal** maps. One characterization of conformal maps, and the reason they are so important to cartographers, is that they preserve angles. Hence angle measurements in the world, *e.g.*, orientation relative to stars, can be used directly on a conformal map. You can see why seafarers might like that... Another characterization, which we will use, is that a conformal map is locally a **scale-rotate**, *i.e.*, a similarity. We begin with maps from $\mathbb{R}^2$ to $\mathbb{R}^2$ before letting a curved surface be the domain.

Suppose we have a map h defined on some path connected, bounded and compact region $\Omega \subset \mathbb{R}^2$ and taking values in $\mathbb{R}^2$, $h\colon \Omega \to \mathbb{R}^2$

$$h(x,y) = \begin{pmatrix} u(x,y) \\ v(x,y) \end{pmatrix}.$$

Let us choose $\{dx, dy\}$ as a basis for 1-forms and $\{e_x, e_y\}$ as the corresponding dual vector basis, then dh can be expressed as the matrix

$$dh = \begin{pmatrix} \partial_x u & \partial_y u \\ \partial_x v & \partial_y v \end{pmatrix}.$$

What condition on the entries of such a 2×2 matrix have to hold for the matrix to represent a similarity? A similarity in $\mathbb{R}^2$ has the form

$$\begin{pmatrix} a & -b \\ b & a \end{pmatrix}$$

where the rotation angle is $\theta = \operatorname{atan2}(b,a)$ and the scale $s = \sqrt{a^2 + b^2}$. Hence dh is a similarity if

$$\partial_x u = \partial_y v \qquad \text{and} \qquad \partial_y u = -\partial_x v.$$

If we treat h as a function from $\mathbb{R}^2$ to $\mathbb{C} \ni u + iv =: h$ we can more concisely write these two conditions as the **Cauchy-Riemann** equations

$$i\partial_x h = \partial_y h.$$

What is the geometric interpretation of this equation? Noting that multiplication by the complex unit i amounts to rotation by $\frac{\pi}{2}$ the Cauchy-Riemann equations say that the derivative along the y-axis agrees with the derivative along the x-axis after it has been rotated a quarter turn counter clockwise. But there is nothing special about the x-axis here. Using the complex valued 1-form dh this amounts to

$$0 = i\,dh(X) - dh(JX) = i\,dh(X) + *dh(X)$$

for X an arbitrary tangent vector and J the linear operator which rotates its argument by $\frac{\pi}{2}$. Here we used the fact that $*dx = dy$ and hence $*dh = -dh \circ J$. J is also known as the complex structure of the domain of our map since $J^2 = -\mathrm{Id}$. We introduce this new symbol J here since we can't perform a $\frac{\pi}{2}$ rotation through multiplication with i if our domain is not $\mathbb{C}$, as is the case for a curved, embedded surface. For such a surface we can choose $J := N\times$, where N is the normal to the surface. Note that $J^2(.) = N \times (N \times (.)) = -\mathrm{Id}(.)$ as required.

Assuming a surface $M \subset \mathbb{R}^3$ with a disk topology, finding a chart amounts to finding a complex valued map $h\colon M \to \mathbb{C}$ which satisfies

$$i\,dh + *dh = 0.$$

While this equation admits an exact solution in the smooth setting, we'll have to settle for an approximate solution in the discretized (piecewise linear) setting of a path connected 2-manifold simplicial complex with boundary. To this end we define the **conformal energy**

$$E_C(h) := \frac{1}{4}\|i\,dh + *dh\|^2$$

and ask for a minimizer

$$\min_h E_C(h)$$

subject to suitable constraints. Since we are now dealing with complex valued 1-forms, $dh \in \Omega^1(M,\mathbb{C})$ we need a **Hermitian** product to replace the usual scalar product.

Exercise 6.1. A Hermitian product is a bi-linear form over (vectors of) complex numbers with the property that

$$\langle u, v\rangle = \overline{\langle v, u\rangle}.$$

Show that for the case of $u, v \in \mathbb{C}$

$$\langle u, v\rangle := \bar{u}v$$

is a Hermitian form and positive definite thus giving us an inner product. Given this Hermitian product show that it gives rise to a Hermitian product for $\mathbb{C}$-valued 1-forms α and β

$$\langle\!\langle \alpha, \beta\rangle\!\rangle = \int_M \bar{\alpha} \wedge *\beta.$$

Now let's work out the form that $E_C(h)$ takes.

Exercise 6.2. Let $h\colon M \to \mathbb{C}$ and assume that its normal derivative along the boundary of M (as before M is an oriented, compact, bounded, path connected 2-manifold with boundary) vanishes. Show that in that case

$$E_C(h) = E_D(h) - A(h)$$

where $E_D(h) = \frac{1}{2}\|dh\|^2$ is the Dirichlet energy of h and A is the area of the image of M

under h

$$A(h) = -\frac{i}{2}\int_M d\bar{h} \wedge dh.$$

Going to a simplicial surface mesh (with disk topology) is now just a matter of the discrete version of $A(h)$, since we already have the discrete version of $E_D(h)$ in the form of the cotan-Laplace matrix.

Exercise 6.3. Show that

$$A(h) = -\frac{i}{4}\sum_{e_{ij}\in E_\partial} \bar{h}_i h_j - \bar{h}_j h_i$$

and conclude that the discrete version of E_C consists of the cotan matrix modified on the boundary. For any boundary edge $e_{ij} \in E_\partial$ give the modification in a form suitable for coding. Hint: use Stokes' theorem.

With the matrices for E_D and A we have a quadratic energy in the simplicial setting of the form

$$\mathrm{E}_C(\mathrm{h}) := \tfrac{1}{2}\bar{\mathrm{h}}\mathrm{C}\mathrm{h}$$

for $\mathrm{C} = \mathrm{L} - 2\mathrm{A}$ with L the cotan-Laplacian, and h the vector of complex coefficients, one to each vertex, representing the 0-form h.

How do we now find the minimizer of the conformal energy? Obviously one minimizer is the constant function $E_C(1) = 0$. So we must exclude that. Additionally we must avoid scaling. To illustrate this point observe that $E_C(ch) = |c|^2 E_C(h)$, *i.e.*, given some function h we can lower its energy simply by scaling with a factor $0 < c < 1$, ultimately returning the zero function as the minimizer. This is easy to avoid by adding the constraint $||h||^2 = 1$.

Exercise 6.4. Consider a positive definite (on all functions orthogonal to the constant function) quadratic form

$$E_C(h) = \tfrac{1}{2}\langle\langle h, Ch\rangle\rangle$$

given by the linear, self-adjoint operator C. Show that the minimizer of $E_C(h)$ subject to $||h||^2 = 1$ and $\langle\langle h, 1\rangle\rangle = 0$ is given by the principal eigenvector h, *i.e.*, the solution to

$$Ch = \lambda h$$

for the smallest non-zero eigenvalue λ. Hint: incorporate the unit norm constraint via a Lagrange multiplier.

In other words all we have to do is solve an eigenvector problem once we have the matrix for our energy. So all we need is a good library to find eigenvectors and hand it our matrix C. As long as we just need the eigenvector associated with the largest or smallest (other than zero) eigenvalue there is a particularly fast way to find it based on power iterations.

Exercise 6.5. Suppose y is the eigenvector associated with the largest eigenvalue Λ of the self-adjoint positive definite linear operator A. Let y^0 be an initial random vector of unit norm with $\langle\langle y^0, y \rangle\rangle \neq 0$ and consider the **power iteration**

$$x = Ay^i \qquad y^{i+1} = \frac{x}{|x|}.$$

Argue that $y^i \to y$ for $i \to \infty$ while $|x|$ converges to Λ. Use this fact to argue that the principal eigenvector, *i.e.*, the eigenvector belonging to the smallest non-zero eigenvalue is found by the inverse power iteration

$$Ax = y^i \qquad y^{i+1} = \frac{x}{|x|}.$$

What happens when A is positive semi-definite as is the case for our matrix C?

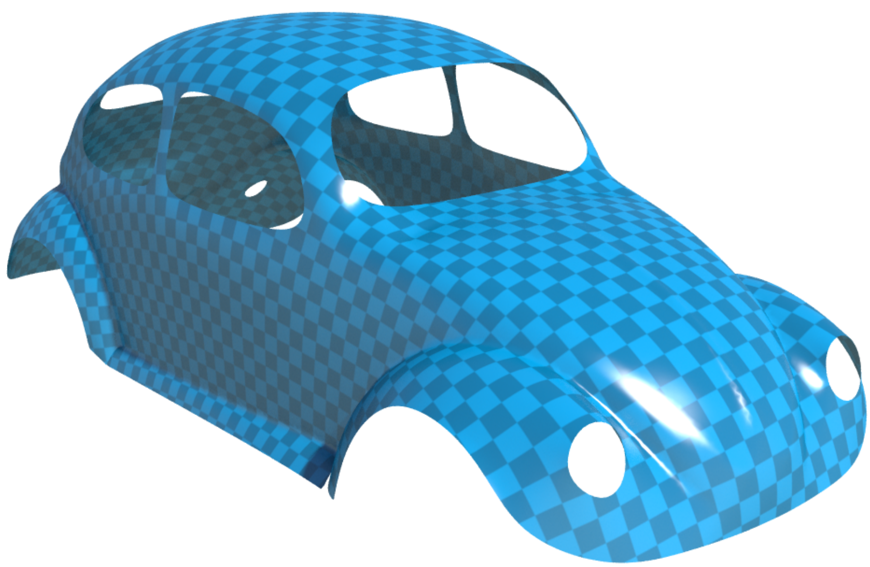

Figure 6.2: *Result of using the first non-zero eigenvector to parameterize the VW bug model. Note that this model is path connected but not simply connected.*

Even though it is seemingly unsophisticated, inverse power iteration can be faster than all other methods. Notice that we have to repeatedly solve problems of the form $Ax = b$ for varying b but constant A. For self-adjoint and full rank A the Cholesky factorization results in an upper triangular matrix U such that $A = U^{\mathsf{T}} U$. While this initial factorization can be as expensive as a full solve, subsequent resolving with a changed right hand side only requires very fast back substitution. Can you think of an easy trick to deal with the fact that our C is not full rank to be able to still use Cholesky factorization?

There is one final detail for implementation.

Exercise 6.6. In the continuous setting, we argued above that the non-zero minimizer of $E_C(h)$, subject to the unit norm constraint $||h||^2 = 1$ solves the principal eigenvalue problem

$$Ch = \lambda h.$$

Show that the discretized version requires solution of the generalized eigenvalue problem

$$\mathrm{Ch} = \lambda \mathrm{Mh}$$

where M is the mass matrix associated with vertices. Show that this is easy to incorporate into the inverse power iteration. Show further that a diagonal M can be incorporated into a modified C such that

$$\tilde{\mathrm{C}}\tilde{\mathrm{h}} = \lambda \tilde{\mathrm{h}},$$

requiring only a method for dealing with a standard eigenvector problem even with $\tilde{\mathrm{C}}$ still being Hermitian. How are h and $\tilde{\mathrm{h}}$ related? Hint: consider $\mathrm{M}^{\frac{1}{2}}$.