The Orientation Manifesto
(for undergraduates)

Timothy E. Goldberg

11 November 2007

An **orientation** of a vector space is represented by an **ordered basis** of the vector space. We think of an orientation as a **twirl**, namely the twirl that rotates the first basis vector to the second, and the second to the third, and so on. Two ordered bases represent the same orientation if they generate the same twirl. (This amounts to the linear transformation taking one basis to the other having positive determinant.) If you think about it carefully, there are only ever two choices of twirls, and hence only two choices of orientation.

Because each $\mathbb{R}^n$ has a standard choice of ordered basis, $\{e_1, e_2, \ldots, e_n\}$ (where e_i is has 1 in the ith coordinates and 0 everywhere else), each $\mathbb{R}^n$ has a standard choice of orientation. The standard orientation of $\mathbb{R}$ is the twirl that points in the positive direction. The standard orientation of $\mathbb{R}^2$ is the counterclockwise twirl, moving from $e_1 = (1, 0)$ to $e_2 = (0, 1)$. The standard orientation of $\mathbb{R}^3$ is a twirl that sweeps from the positive x direction to the positive y direction, and up the positive z direction. It's like a directed helix, pointed up and spinning in the counterclockwise direction if viewed from above. See Figure 1.

An orientation of a curve, or a surface, or a solid body, is really a choice of orientations of every single tangent space, in such a way that the twirls all agree with each other. (This can be made horribly precise, when necessary.)

There are several ways for a manifold to pick up an orientation.

(1) **From the surrounding space.** If the manifold is n-dimensional and sits in $\mathbb{R}^n$, then the manifold can just pick up the standard orientation of $\mathbb{R}^n$. Examples of this are a curve living in the line $\mathbb{R}$, or a surface living in the plane $\mathbb{R}^2$, or a solid body living in the space $\mathbb{R}^3$.

(2) **From the parametrization.** If the manifold is the image of a function, then you can give the manifold the orientation that comes from the standard

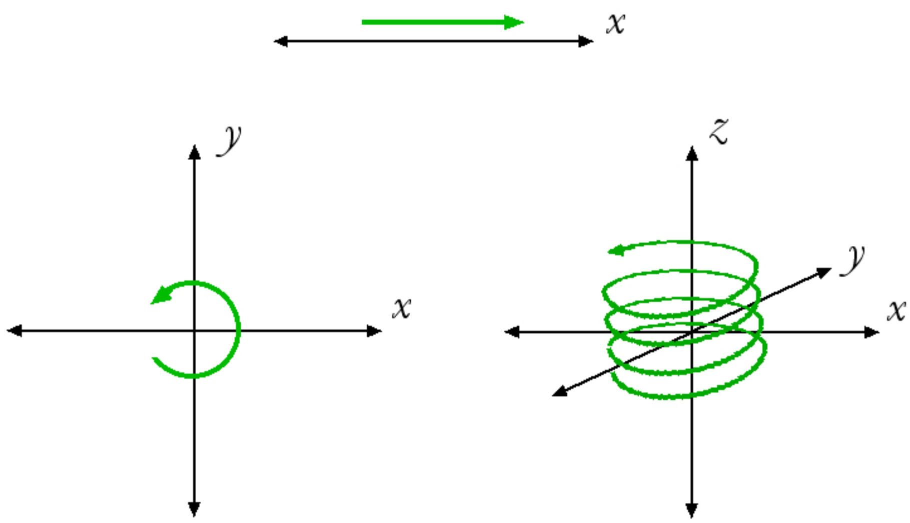

Figure 1: The standard orientations of $\mathbb{R}$, $\mathbb{R}^2$, and $\mathbb{R}^3$.

orientation of the function's domain. For example, suppose S is a surface that is given by the image of the map $k\colon \mathbb{R}^2 \to \mathbb{R}^3, (u, v) \mapsto k(u, v)$. Then at each $p \in \mathbb{R}^2$, the **total derivative** of k at p is a linear map, $dk_p\colon \mathbb{R}^2 \to \mathbb{R}^3$. Really, we should think of dk_p as a linear map from the tangent plane of $\mathbb{R}^2$ at p to the tangent plane of S at $k(p)$. Since the standard basis $e_1 = (1, 0), e_2 = (0, 1)$ of $\mathbb{R}^2$ generates the twirl that is the standard orientation of $\mathbb{R}^2$, it makes sense to give the tangent plane to S at $k(p)$ the twirl, or orientation, given by the ordered basis $\{dk_p(e_1), dk_p(e_2)\}$. See Figure 2. Of course, we can also compute these vectors with the formulas

$$dk_p(e_1) = \frac{\partial k}{\partial u}(p) \quad \text{and} \quad dk_p(e_2) = \frac{\partial k}{\partial v}(p).$$

So if you parametrize a surface, or indeed any manifold, and you're unsure which orientation you've given it with that particular parametrization, you can compute these tangent vectors and think about what twirl they represent.

(3) **From picking a side.** (This only applies to an n-dimensional manifold sitting in $\mathbb{R}^{n+1}$, like a surface sitting in $\mathbb{R}^3$.) Picking a side of a manifold means choosing a continuous normal vector field on the manifold. we think of the normal vectors as pointing at the side we picked. If you can choose one of these, you can choose another one by having each vector point in the opposite

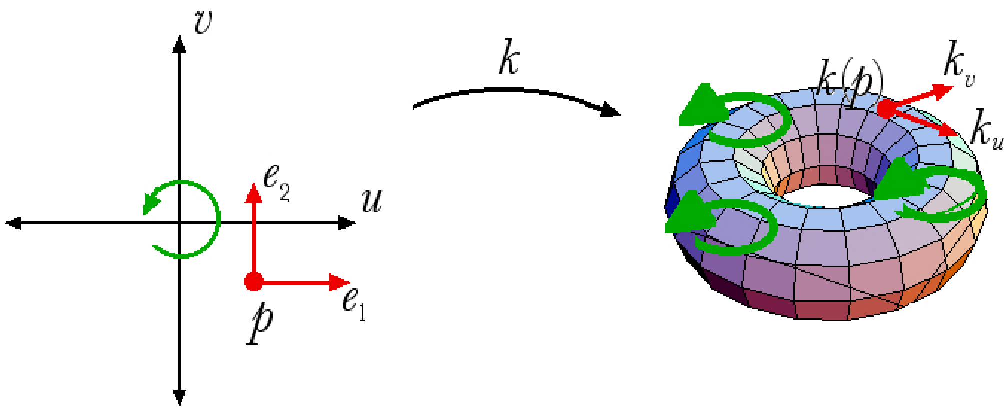

Figure 2: Orientation induced by parametrization. Here, $k_u = \frac{\partial k}{\partial u}$ and $k_v = \frac{\partial k}{\partial v}$.

direction. In fact, a manifold is **orientable** exactly if you can choose a normal vector field, i.e. exactly if it has two different sides. The Möbius strip is the standard example of a nonorientable manifold, because it only has one side. You cannot find a continuous normal vector field on the Möbius strip.

So suppose we pick a side of the manifold, meaning a continuous normal vector field. There are two ways to orient the manifold, two choices of twirl. The choice of twirl, or orientation, induced by our choice of side is the one which, if we lift the twirl in the direction of the normal vector, we obtain the standard twirl of the vector space $\mathbb{R}^{n+1}$ in which it sits. For a surface S in $\mathbb{R}^3$ and a choice of normal vector field, the induced representation on S is the one which, when viewed from the arrow of one of the normal vectors, is twirling counterclockwise. If we look at the surface from the other side, this orientation will be the one twirling clockwise. Imagine an inner tube floating on the surface of a swimming pool, and imagine someone sets it spinning. Whichever direction it looks to be spinning from out of the water, from under the water it will be looking to be spinning the other direction.

(4) **From being a boundary.** If your manifold is actually the boundary ∂M of another manifold M, and if M is oriented, then there is a standard way that ∂M inherits an orientation. If M is orientable, then so is its boundary ∂M, so there are two choices of continuous normal vector fields on ∂M. One of them will point toward M (or inside), and the other will point away from M (or

outside). For the inherited orientation on $\partial \mathsf{M}$, we choose the normal vectors that point to the outside, and then take the orientation that this induces on $\partial \mathsf{M}$, as in (3) above.

For a surface S which is equal to the boundary of an oriented solid M, I like to imagine the induced orientation as a little boat. The boat has a mast, and we know which direction the boat is pointing. The solid M is the water, and the surface S is the surface of the water, and the boat floats on S. The mast points in the normal direction to the outside of M. The induced orientation on S is always the one which spins our boat to the left. See Figure 3.

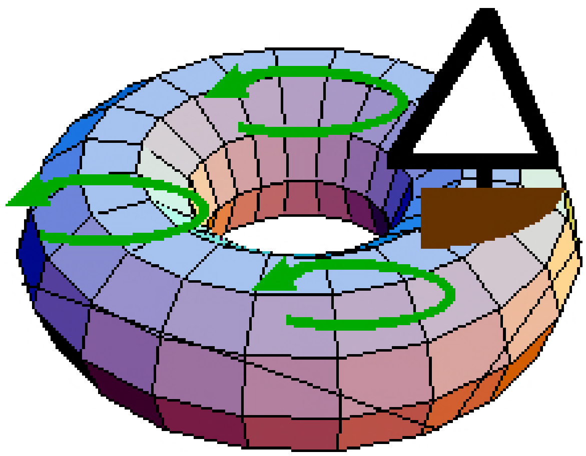

Figure 3: A little boat on a torus, at the mercy of a fierce counterclockwise orientation. Note that the mast points in the outward normal direction to the surface.

For a curve C which is equal to the boundary of an oriented solid S, I like to imagine that S is a swimming pool and C is the boundary of the pool. We can find the induced orientation, or direction, of C by imagining in what direction we would need to move around the pool to generate the given twirl of the water in the pool.

Notice that there's no particular reason that we use the outward pointing normal vectors to induce orientation on boundaries. We could just as easily use the inward pointing normal. (Although in this case, Stokes' theorem would need an extra minus sign, wouldn't it?) I suspect that we use the outward pointing normal because we like to imagine that we are viewing the surface from the outside. Or maybe all of us like to imagine little boats floating on the surface.

From the discussion above, we can see why if we are going to glue two oriented manifolds M_1 and M_2 together along their common boundary to form a third manifold, then M_1 and M_2 must induce the opposite orientations on their boundaries. See Figure 4.

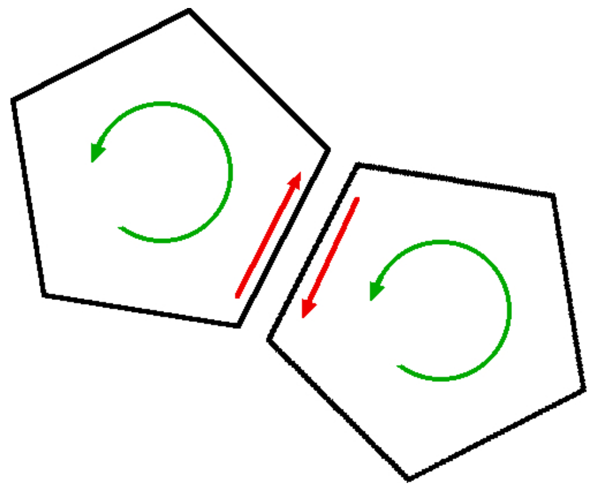

Figure 4: Opposite orientations attract.

Here are the main points regarding integrals.

- You can integrate a function $\mathbb{R}^n \to \mathbb{R}$ over a region of $\mathbb{R}^n$. Here a region of $\mathbb{R}^n$ essentially means an n-dimensional manifold, or the union of some n-dimensional manifolds, in $\mathbb{R}^n$. Essentially, orientations don't come into it. These are the usual integrals you were familiar with before this class.

- You can integrate a differential form of degree n over an *oriented* n-dimensional manifold, or over the union of some oriented n-dimensional manifolds. If you switch the given orientation on the manifold to its opposite, the value of the integral is negated.

- Integrals of functions and integrals of differential forms are not unrelated. Given a region R in $\mathbb{R}^n$, give it the orientation induced by the standard orientation of $\mathbb{R}^n$. Given a function $f\colon R \to \mathbb{R}$ defined on R, form the differential n-form $\alpha = f\, dx_1 \ldots dx_n$. Then the integral of f over $\mathbb{R}$ equals the integral of α over R with this orientation.