Math 3204: Calculus IV[1] Mikhail Lavrov

Lecture 28: Putting it all together

December 4, 2023 *Kennesaw State University*

0 Why?

Today's lecture notes turned out to be really, really long.

I will not cover most of this in class. My goal for the last lecture is just to show you how the different theorems we've learned this semester fit together, and if there's time at the end, we can spend it on review.

But I wanted to at least present, as a resource for you to refer to if you'd like, the way that everything we've done generalizes to k-dimensional objects in n-dimensional space. I think that this is the only way to see how the definitions we've made are the *right* definitions to make, and the integrals we've taken are the *right* integrals to take.

1 Vector integral theorems

Here is a summary of the theorems we've seen that relate two vector integrals over objects of different dimensions.

Theorem 1.1 (Fundamental theorem of line integrals)**.** *Let C be an oriented curve in $\mathbb{R}^3$ that starts at point $\mathbf{a}$ and ends at point $\mathbf{b}$; let $f\colon \mathbb{R}^3 \to \mathbb{R}$ be a scalar function. Then*

$$\int_C \nabla f \cdot \mathbf{T}\,ds = f(\mathbf{b}) - f(\mathbf{a}).$$

As a special case, if C is a closed curve (it starts where it ends), then the line integral of ∇f around C is 0.

Theorem 1.2 (Stokes' theorem)**.** *Let S be an oriented surface in $\mathbb{R}^3$ and let C be its compatibly oriented boundary; let $\mathbf{F}\colon \mathbb{R}^3 \to \mathbb{R}^3$ be a vector field. Then*

$$\iint_S (\nabla \times \mathbf{F}) \cdot \mathbf{n}\,dS = \int_C \mathbf{F} \cdot \mathbf{T}\,ds.$$

As a special case, if S is a closed surface (it has no boundary), then the flux integral of $\nabla \times \mathbf{F}$ across S is 0.

Green's theorem is another special case of this theorem—it is what we get if S is a flat region in $\mathbb{R}^2$.

Theorem 1.3 (Divergence theorem)**.** *Let D be a solid region in $\mathbb{R}^3$ and let S be its outward-oriented boundary; let $\mathbf{F}\colon \mathbb{R}^3 \to \mathbb{R}^3$ be a vector field. Then*

$$\iiint_D \nabla \cdot \mathbf{F}\,dV = \iint_S \mathbf{F} \cdot \mathbf{n}\,dS.$$

[1]This document comes from the Math 3204 course webpage: http://facultyweb.kennesaw.edu/mlavrov/courses/3204-fall-2023.php

It is easy to superficially say that these theorems are all saying the same thing: integrating some kind of "vector-like derivative" of a function over an object gives the same result as integrating the original function over the boundary of that object. I find this kind of observation somewhat frustrating, because it doesn't tell you enough about what to expect in situations you *haven't* seen before. What tells us that the curl $\boldsymbol{\nabla} \times \mathbf{F}$ is the correct derivative to take in Stokes' theorem, and the divergence $\boldsymbol{\nabla} \cdot \mathbf{F}$ is the correct derivative to take in the divergence theorem?[2] How do we know that the flux integral across a surface is the correct kind of double integral to use in these theorems?

2 Manifolds

First, we need to say a little bit about the common generalization of a point, curve, or surface in $\mathbb{R}^3$: this is called a **manifold**.

2.1 Defining manifolds

If you are the sort of person who is happy with intuitive geometric feelings, you might be happy just saying that a curve is a 1-dimensional object, and a surface is a 2-dimensional object. If you prefer rigorous definitions, you might have the question: what makes a manifold k-dimensional? There are two answers:

- On a theoretical level that we cannot make rigorous in *this* course, the answer is that a k-dimensional manifold is something that "locally looks like a piece of $\mathbb{R}^k$". For example, a surface is 2-dimensional, because locally it looks like a piece of the plane.
- On a practical level that is useful for our purposes, the answer is that a k-dimensional manifold is something that can be given a k-dimensional parameterization: it is a mostly-injective image of a function $\mathbf{r}(u_1, u_2, \ldots, u_k)$, where u_i ranges over an interval $[a_i, b_i]$ for every i.

In our case, a 1-dimensional manifold in $\mathbb{R}^3$ is a curve: something with a parameterization $\mathbf{r}(t)$, where $t \in [a, b]$. A 2-dimensional manifold in $\mathbb{R}^3$ is a surface: something with a parameterization $\mathbf{r}(u, v)$, where $(u, v) \in [a, b] \times [c, d]$. A 3-dimensional manifold in $\mathbb{R}^3$: is a solid region: usually, we describe these just in terms of x, y, and z (*three* variables!) but in principle, we could use a uvw-substitution to describe them, which is like using a 3-variable parameterization $\mathbf{r}(u, v, w)$. A 0-dimensional manifold is a point: it takes no variables to describe.

We are also happy thinking about k-dimensional manifolds that have several pieces, each with its own k-dimensional parameterization.

2.2 Oriented manifolds

All of these objects can be **oriented**, which means different things depending on context. For a curve, orientation means giving it a starting point and an ending point, or a direction of travel along the curve. For a surface, orientation means labeling one of the two sides of the curve as an "inside"

[2] Well, the fact that it's called the "divergence theorem" may be a subtle hint in that direction, but if we had never seen that theorem before, presumably we wouldn't know what it's called, either.

and an "outside". The extreme cases—orientations of points and solid regions—are a bit silly in $\mathbb{R}^3$, and we can just think of those cases as having a "positive" or "negative" orientation.

Importantly, all of these orientations (except maybe the point) can be inferred from the parameterization we gave our manifold:

- For a curve, our parameterization $\mathbf{r}(t)$ has a derivative $\frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t}$. This gives a tangent vector to the curve at the point $\mathbf{r}(t)$, which tells us a direction of travel along the curve.
- For a surface, our parameterization $\mathbf{r}(u, v)$ has two partial derivatives $\frac{\partial \mathbf{r}}{\partial u}$ and $\frac{\partial \mathbf{r}}{\partial v}$. Their cross product gives a normal vector to the curve at the point $\mathbf{r}(u, v)$, which we interpret as pointing from "inside" to "outside".
- For a solid region, if we have a uvw-substitution and interpret it as a 3-variable parameterization $\mathbf{r}(u, v, w)$, then we can compute the Jacobian determinant $\frac{\partial(x,y,z)}{\partial(u,v,w)}$ of that substitution; the sign of the Jacobian determinant tells us if the orientation is positive or negative.

It looks like we need specialized knowledge related to curves and surfaces to interpret how, exactly, the parameterization gives us an orientation. This is not exactly true: we only need this specialized knowledge if we want to put the orientation in a human-legible form. In the abstract context of considering k-dimensional manifolds in $\mathbb{R}^n$, it is enough to know that there's *some* notion of orientation that's encoded in the parameterization.

2.3 Boundaries of manifolds

For $k \geq 1$, a k-dimensional manifold has a boundary, which is a $(k-1)$-dimensional manifold. The boundary of a curve in $\mathbb{R}^3$ consists of the endpoints of that curve. The boundary of a surface in $\mathbb{R}^3$ is the curve that goes around the edge of the surface. The boundary of a solid region in $\mathbb{R}^3$ is the surface enclosing that region.

It is possible for the boundary to be empty: for example, a curve could be a closed curve. We think of it as starting where it ends, and we give it a parameterization $\mathbf{r}(t)$ with $t \in [a, b]$ such that $\mathbf{r}(a) = \mathbf{r}(b)$; however, that doesn't mean that we think of $\mathbf{r}(a)$ as a boundary point of the curve! Geometrically, it's just a point on the curve like any other.

The boundary of an oriented manifold is also oriented; moreover, there's a notion of compatibility between the orientation of a manifold and the orientation of its boundary. For example, if a curve starts at $\mathbf{r}(a)$ and ends at $\mathbf{r}(b)$, the compatible orientation of these points makes $\mathbf{r}(b)$ a "positive" point and $\mathbf{r}(a)$ a "negative" point. For a positively oriented solid region, the compatible orientation of its boundary is the outward orientation.

In the case of surfaces and their boundaries, we had a completely algebraic way of finding boundaries based on the parameterization: see Lecture 23. This is conceptually important, because that method gave us *oriented* boundaries, and so it gave us an algebraic interpretation of what it means for boundaries to be compatibly oriented. The method we used in Lecture 23 for surfaces generalizes to arbitrary dimensions (of both manifolds and the spaces they live in).

In this lecture, we wil introduce some new notation: we will write ∂M for the boundary of the manifold M.

3 Differential forms

Differential forms are what we need to talk about integration over manifolds. The machinery they bring in is admittedly more complicated than what we need to understand vector integrals. However, there are no *arbitrary* secrets specific to a particular situation to learn about differential forms, and it is possible to state a single theorem that unites all three of the theorems we began this lecture with.

3.1 Building differential forms

The fundamental building blocks of differential forms on $\mathbb{R}^3$ are the coordinate differentials $\mathrm{d}x$, $\mathrm{d}y$, and $\mathrm{d}z$. (Working in full generality over $\mathbb{R}^n$, we would have n different coordinate differentials $\mathrm{d}x_1, \mathrm{d}x_2, \ldots, \mathrm{d}x_n$.) They are combined using the "exterior" or "wedge" product written using the symbol $\wedge$. A single differential $\mathrm{d}x$, $\mathrm{d}y$, or $\mathrm{d}z$ intuitively represents an infinitesimal oriented length; a wedge product like $\mathrm{d}x \wedge \mathrm{d}y$, $\mathrm{d}y \wedge \mathrm{d}z$, or $\mathrm{d}z \wedge \mathrm{d}x$ intuitively represents an infinitesimal oriented area; a triple wedge product like $\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z$ intuitively represents an infinitesimal oriented volume.

All of these can be multiplied by arbitrary scalar functions and added together. However, we prefer not to combine terms with a different number of differentials; it is generally not useful to consider a sum like $x^2\,\mathrm{d}y - \mathrm{d}x \wedge \mathrm{d}z$. We say that a **differential k-form** is one which only includes k-fold wedge products of the coordinate differentials. For example, a generic differential 1-form is something that can be written as

$$M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z$$

for some scalar functions M, N, and P.

We can define the wedge product of *any* two differential forms. The wedge product is associative, distributes over addition, and commutes with scalar multiplication. When simplifying wedge products of differentials, we obey two rules:

1. For any two variables u and v, $\mathrm{d}u \wedge \mathrm{d}v = -(\mathrm{d}v \wedge \mathrm{d}u)$.
2. For any variable u, $\mathrm{d}u \wedge \mathrm{d}u = 0$.

This wedge product is actually a common generalization of several useful vector operations. For example, the wedge product of a 1-form and a 2-form

$$(M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z) \wedge (U\,\mathrm{d}y \wedge \mathrm{d}z + V\,\mathrm{d}z \wedge \mathrm{d}x + W\,\mathrm{d}x \wedge \mathrm{d}y)$$

will simplify to the 3-form $(MU + NV + PW)\,\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z$, so it is analogous to the dot product of vectors $M\,\mathbf{i} + N\,\mathbf{j} + P\,\mathbf{k}$ and $U\,\mathbf{i} + V\,\mathbf{j} + W\,\mathbf{k}$. The wedge product of two 1-forms will simplify to a 2-form, and it corresponds to the cross product of two vectors. The wedge product of three 1-forms will simplify to a 3-form, and it corresponds to the determinant of a 3×3 matrix.

(All of these examples are specific to differential forms on $\mathbb{R}^3$, but that doesn't mean that the wedge product doesn't do anything interesting on $\mathbb{R}^n$ for other values of n—it just has other, less familiar meanings in those cases.)

3.2 Exterior derivatives

There is another abstract operation we can define on differential forms: the exterior derivative. This is an operation d that turns 0-forms (or functions) into 1-forms, 1-forms into 2-forms, and so forth.

The rule is:

- For a 0-form f in $\mathbb{R}^3$, its exterior derivative $\mathrm{d}f$ is defined by $\mathrm{d}f = \frac{\partial f}{\partial x}\,\mathrm{d}x + \frac{\partial f}{\partial y}\,\mathrm{d}y + \frac{\partial f}{\partial z}\,\mathrm{d}z$. This generalizes to n dimensions if you like: the formula becomes $\mathrm{d}f = \frac{\partial f}{\partial x_1}\,\mathrm{d}x_1 + \frac{\partial f}{\partial x_2}\,\mathrm{d}x_2 + \cdots + \frac{\partial f}{\partial x_n}\,\mathrm{d}x_n$.
- To take the exterior derivative of a k-form with $k \geq 1$, we apply the rule above to every scalar function that appears in the k-form; multiplication by that scalar function becomes a wedge product with its exterior derivative. For example,

  $$\mathrm{d}(M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z) = \mathrm{d}M \wedge \mathrm{d}x + \mathrm{d}N \wedge \mathrm{d}y + \mathrm{d}P \wedge \mathrm{d}z$$

  is the rule for the exterior derivative of a 1-form in $\mathbb{R}^3$.

  We simplify this expression using the already-mentioned rules for simplifying wedge products.

In $\mathbb{R}^3$, the exterior derivative is the common generalization of gradient, curl, and divergence. For a 0-form f, its exterior derivative $\mathrm{d}f$ (shown above) corresponds to the gradient of f. For a 1-form $M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z$, its exterior derivative will end up simplifying to

$$\mathrm{d}(M\,\mathrm{d}x + N\,\mathrm{d}y + P\,\mathrm{d}z) = \left(\frac{\partial P}{\partial y} - \frac{\partial N}{\partial z}\right)\mathrm{d}y \wedge \mathrm{d}z + \left(\frac{\partial M}{\partial z} - \frac{\partial P}{\partial x}\right)\mathrm{d}z \wedge \mathrm{d}x + \left(\frac{\partial N}{\partial x} - \frac{\partial M}{\partial y}\right)\mathrm{d}x \wedge \mathrm{d}y$$

which corresponds to the curl of $\mathbf{F} = M\,\mathbf{i} + N\,\mathbf{j} + P\,\mathbf{k}$. For a 2-form $U\,\mathrm{d}y \wedge \mathrm{d}z + V\,\mathrm{d}z \wedge \mathrm{d}x + W\,\mathrm{d}x \wedge \mathrm{d}y$, its exterior derivative will end up simplifying to $\left(\frac{\partial U}{\partial x} + \frac{\partial V}{\partial y} + \frac{\partial W}{\partial z}\right)\mathrm{d}x \wedge \mathrm{d}y \wedge \mathrm{d}z$, which corresponds to the divergence of $\mathbf{G} = U\,\mathbf{i} + V\,\mathbf{j} + W\,\mathbf{k}$.

An important property of the exterior derivative, which we proved in an earlier lecture, is that for any differential form ω, $\mathrm{d}(\mathrm{d}\omega) = 0$. We have the diagram

$$0\text{-forms} \xrightarrow{\mathrm{d}} 1\text{-forms} \xrightarrow{\mathrm{d}} 2\text{-forms} \xrightarrow{\mathrm{d}} 3\text{-forms}$$

and if we start anywhere, and apply the arrow twice, we get 0. In fancy algebraic terminology, this makes the diagram a "chain complex".

3.3 Integrating differential forms

Messing around with differential forms is useless until we can relate them to something we already understand—and statements like "dx intuitively represents an infinitesimal oriented length" don't mean anything until we give a way to interpret dx non-intuitively.

In fact, the proper way to understand what differential forms *are* is this: a differential k-form represents something we can integrate over a k-dimensional oriented manifold.

If you have a differential k-form, and a k-dimensional manifold, then you can express our coordinate differentials in terms of coordinates *on the manifold*: differentials of the parameterization variables!

Once we've done that, our integral of the differential form over the manifold turns into a standard k-fold iterated integral with respect to the variables of the parameterization.

For example, if you have a curve in $\mathbb{R}^3$ described by a parameterization $\mathbf{r}(t)$ and a differential form $\omega = M\,\mathrm{d}x + N\,\mathrm{d}y + N\,\mathrm{d}t$, you can use the exterior derivative to express $\mathrm{d}x$, $\mathrm{d}y$, and $\mathrm{d}z$ all in terms of $\mathrm{d}t$ (which will involve differentiating the components of $\mathbf{r}(t) = (x(t), y(t), z(t))$). This puts ω into the form $f(t)\,\mathrm{d}t$ for some function $f(t)$. Now, integrating ω is as simple as integrating $f(t)$ with respect to t. The bounds on t are whatever they were in the parameterization $\mathbf{r}(t)$.

In $\mathbb{R}^3$, this description will tell us that the integral of a 1-form is a line integral of the corresponding vector field, and the integral of a 2-form is a flux integral of the corresponding vector field—without knowing in advance that there's anything special about line integrals and flux integrals.

4 The generalized Stokes' theorem

We now have all the vocabulary we need to state a single theorem uniting all of our theorems. This theorem is called the generalized Stokes' theorem. It says:

Theorem 4.1. *Let M be a $(k+1)$-dimensional oriented manifold in $\mathbb{R}^n$ (where $0 \le k \le n-1$), and let ω be a differential k-form defined on M. Then*

$$\int_M \mathrm{d}\omega = \int_{\partial M} \omega.$$

(Here, as mentioned earlier, ∂M denotes the boundary of M; it must be oriented compatibly with the orientation of M.)

We will not prove this theorem, but the proof of the divergence theorem in the previous lecture captures the spirit of what's necessary to prove it. The rest is just mastering the notation of differential forms.

One thing worth noting is that we said "let ω be a differential k-form defined on M". This is somewhat more general than the way we've been using theorems like this in practice; we've written our integrands in terms of x, y, and z, which meant that our differential forms were defined on all of $\mathbb{R}^3$.

However, earlier in this lecture, when we defined the integral of ω on M, our first step was using the parameterization of M to write ω in terms of the parameterization variables. In principle, we could skip that step and just express ω in terms of those variables to begin with—then, we don't need ω to be defined anywhere outside M in the first place.

More commonly, we take advantage of this "defined on M" clause to apply the theorem to differential forms that have poles of some kind—as long as those poles are not contained in M. For example, if we divide by $x^2 + y^2 + z^2$ in the definition of ω, that's fine—as long as M does not contain the point $(0, 0, 0)$.

One of the theorems at the beginning of these notes seems particularly unlike the the generalized Stokes' theorem. The fundamental theorem of line integrals doesn't even have integrals on both sides! What's up with that?

To make sense of this, we must define what it means to integrate over a point. A point is a 0-dimensional manifold. A 0-form (the thing we can integrate over a 0-dimensional manifold) is a function. What is the only thing we can do with a function and a point? Evaluate that function at that point! So the integral of a function $f\colon \mathbb{R}^n \to \mathbb{R}$ over a point $\mathbf{p} \in \mathbb{R}^n$ is simply the value $f(\mathbf{p})$.

An *oriented* point is just a point that can be either positive or negative. The integral of f over a positive point $\mathbf{p}$ is $f(\mathbf{p})$, and the integral of f over a negative point $\mathbf{p}$ is $-f(\mathbf{p})$. How does this come up? Well, if we have a curve C that starts at point $\mathbf{a}$ and ends at point $\mathbf{b}$, its compatibly-oriented boundary consists of the *positively-oriented* point $\mathbf{b}$ and the *negatively-oriented*point $\mathbf{a}$, so the integral of f over its boundary is precisely $f(\mathbf{b}) - f(\mathbf{a})$.

4.1 Boundary of a boundary

There is a neat geometrical consequence of the generalized Stokes' theorem. Let ω be a k-form and let M be a $(k+2)$-dimensional oriented manifold. Then by applying the generalized Stokes' theorem twice, we get

$$\int_M \mathrm{d}(\mathrm{d}\omega) = \int_{\partial M} \mathrm{d}\omega = \int_{\partial(\partial M)} \omega$$

where $\partial(\partial M)$ denotes the boundary of the boundary of M. The first integral here is always 0, because $\mathrm{d}(\mathrm{d}\omega) = 0$ for any differential form ω. Therefore the other two integrals are also 0. In particular, the third integral (which is a integral of an *arbitrary* k-form over $\partial(\partial M)$) is always 0.

How can this be? If $\partial(\partial M)$ were any nontrivial manifold, then surely there would be some k-forms that gave us a nonzero integral. The only possibility is that $\partial(\partial M)$ is guaranteed to be empty: the boundary of a manifold has no boundary of its own.

5 Scalar integrals

A scalar integral is an integral of a scalar function $f\colon \mathbb{R}^n \to \mathbb{R}$ over a manifold M that does *not* care about the orientation of M. These have not been the focus of our attention for a few lectures, but we would also like to understand them in general.

The general approach here[3] is an extension of the Jacobian determinant. Suppose that $\mathbb{R}^n$ has variables $x_1, x_2, \ldots, x_n$, and that M is a k-dimensional manifold on which these are parameterized by $u_1, u_2, \ldots, u_k$. We define the matrix J to be the matrix of partial derivatives

$$J = \begin{bmatrix} \frac{\partial x_1}{\partial u_1} & \frac{\partial x_1}{\partial u_2} & \cdots & \frac{\partial x_1}{\partial u_k} \\ \frac{\partial x_2}{\partial u_1} & \frac{\partial x_2}{\partial u_2} & \cdots & \frac{\partial x_2}{\partial u_k} \\ \frac{\partial x_3}{\partial u_1} & \frac{\partial x_3}{\partial u_2} & \cdots & \frac{\partial x_3}{\partial u_k} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial x_n}{\partial u_1} & \frac{\partial x_n}{\partial u_2} & \cdots & \frac{\partial x_n}{\partial u_k} \end{bmatrix}.$$

[3] Which I got a nicer explanation of by asking a question on Math StackExchange: https://math.stackexchange.com/questions/4819542/scalar-integrals-in-higher-dimensions.

What have we already done with this matrix?

- When performing a uv-substitution in $\mathbb{R}^2$ or a uvw-substitution in $\mathbb{R}^3$, we use the Jacobian determinant $\det(J)$ as a scaling factor. In oriented integrals, we left it as it is; for scalar integrals, we multiply by $|\det(J)|$, its absolute value.

  (This cannot be the general rule for scalar integrals of all dimensions, by the way: when the manifold M has lower dimension than $\mathbb{R}^n$, J is not a square matrix, and has no determinant.)

- When we give a curve C in $\mathbb{R}^n$ a parameterization $\mathbf{r}(t)$, the matrix J is just a column vector: the vector $\frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t}$. In a scalar integral over C, we multiply by $\left\|\frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t}\right\|$: the sum of squares of the entries of J.

- The factor $\left\|\frac{\partial \mathbf{r}}{\partial u} \times \frac{\partial \mathbf{r}}{\partial v}\right\|$ for scalar surface integrals also seems to involve the matrix J somehow, but it might not be immediately obvious how to write that relationship down in a formula.

5.1 The Gram determinant

The common generalization of all of these is the following rule. When integrating over *any* k-dimensional object in *any* number of dimensions, our "k-dimensional volume element" (which is written $\mathrm{d}s$, or $\mathrm{d}A$, or $\mathrm{d}V$ in the special cases we've seen) can be expressed in terms of our parameters $u_1, \ldots, u_k$ as

$$\sqrt{\det(J^\mathsf{T} J)}\,\mathrm{d}u_1\,\mathrm{d}u_2 \ldots \mathrm{d}u_k.$$

Here, J^T is the transpose of the matrix J; we take the matrix product of the $k \times n$ matrix J^T and the $n \times k$ matrix J to get a $k \times k$ result, and then we take its determinant. (That determinant is guaranteed to be positive, so we can always take its square root.)

If you want some terminology: the matrix $J^\mathsf{T} J$ is called the **Gram matrix** of the columns of J, and its determinant is called the **Gram determinant**. It is a more general concept than what we're using here. In general, $\sqrt{\det(J^\mathsf{T} J)}$ computes the k-dimensional volume of the parallelepiped[4] whose sides are parallel and congruent to the columns of J. In our case, the way we've defined J, the columns of J are the partial derivatives $\frac{\partial \mathbf{r}}{\partial u_1}$ through $\frac{\partial \mathbf{r}}{\partial u_k}$.

The Gram matrix $J^\mathsf{T} J$ can also be defined in a different, equivalent way: its (i, j) entry is the dot product of the i^{th} column of J and the j^{th} column.

Now, if we ever need to, we could integrate a function over a 5-dimensional manifold in 17-dimensional space. (It won't be fun, but we know how to do all of the steps!)

Let's compare the Gram determinant approach to several situations in $\mathbb{R}^3$:

- If $k = 1$, the manifold M is a curve with parameterization $\mathbf{r}(t)$, and J has a single column vector $\frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t}$. Computing $J^\mathsf{T} J$ means taking the dot product of $\frac{\mathrm{d}\mathbf{r}}{\mathrm{d}t}$ with itself, which is exactly the same as adding up the squares of its entries.

- If $k = 2$, then the $\sqrt{\det(J^\mathsf{T} J)}$ formula is equivalent to the $\left\|\frac{\partial \mathbf{r}}{\partial u} \times \frac{\partial \mathbf{r}}{\partial v}\right\|$ after simplification, but this is not obvious. As we'll see in an example, the $\sqrt{\det(J^\mathsf{T} J)}$ formula can sometimes be easier to use!

[4] A hard-to-spell word that generalizes "parallelogram" to more than 2 dimensions.

- If $k = 3$, then J is a square matrix, and has a determinant; we can write

$$\sqrt{\det(J^\mathsf{T} J)} = \sqrt{\det(J^\mathsf{T}) \cdot \det(J)} = \sqrt{\det(J)^2} = |\det(J)|.$$

  The formula via Gram determinant is a bit more unwieldy than just using $|\det(J)|$, but that's the price we pay for generality.

5.2 An example

Out of an excess of bravery, let's try to compute $\sqrt{\det(J^\mathsf{T} J)}$ for the spherical parameterization $\mathbf{r}(\theta, \phi) = (\cos\theta \sin\phi, \sin\theta \sin\phi, \cos\phi)$. We'll take the dot-product interpretation of $J^\mathsf{T} J$:

$$J^\mathsf{T} J = \begin{bmatrix} \frac{\partial \mathbf{r}}{\partial \theta} \cdot \frac{\partial \mathbf{r}}{\partial \theta} & \frac{\partial \mathbf{r}}{\partial \theta} \cdot \frac{\partial \mathbf{r}}{\partial \phi} \\ \frac{\partial \mathbf{r}}{\partial \phi} \cdot \frac{\partial \mathbf{r}}{\partial \theta} & \frac{\partial \mathbf{r}}{\partial \phi} \cdot \frac{\partial \mathbf{r}}{\partial \phi} \end{bmatrix}$$

where $\frac{\partial \mathbf{r}}{\partial \theta} = -\sin\theta \sin\phi\, \mathbf{i} + \cos\theta \sin\phi\, \mathbf{j}$ and $\frac{\partial \mathbf{r}}{\partial \phi} = \cos\theta \cos\phi\, \mathbf{i} + \sin\theta \cos\phi\, \mathbf{j} - \sin\phi\, \mathbf{k}$.

The $(1,1)$ entry is

$$\frac{\partial \mathbf{r}}{\partial \theta} \cdot \frac{\partial \mathbf{r}}{\partial \theta} = \sin^2\theta \sin^2\phi + \cos^2\theta \sin^2\phi,$$

which we simplify to $\sin^2\phi$ by factoring out $\sin^2\theta + \cos^2\theta = 1$. The $(1,2)$ and $(2,1)$ entries are both equal to

$$\frac{\partial \mathbf{r}}{\partial \theta} \cdot \frac{\partial \mathbf{r}}{\partial \phi} = \cos\theta \cos\phi(-\sin\theta \sin\phi) + \sin\theta \cos\phi \cos\theta \sin\phi + 0$$

whih simplifies to 0. Finally, the $(2,2)$ entry is

$$\frac{\partial \mathbf{r}}{\partial \phi} \cdot \frac{\partial \mathbf{r}}{\partial \phi} = \cos^2\theta \cos^2\phi + \sin^2\theta \cos^2\phi + \sin^2\phi$$

which simplifies twice: the first two terms combine into $\cos^2\phi$, and then this combines with the last term to just get 1. Therefore

$$J^\mathsf{T} J = \begin{bmatrix} \sin^2\phi & 0 \\ 0 & 1 \end{bmatrix}.$$

I feel like this is much less painful than the method using cross products! It reveals a big part of the reason for the simplification: the two partial derivatives are orthogonal to each other, which is why we like this parameterization to begin with.

For surface integrals with this parameterization, we will use $\sqrt{\det(J^\mathsf{T} J)} = \sqrt{\sin^2\phi} = |\sin\phi|$. Usually, the domain of ϕ is $[0, \pi]$, in which case, we can just write $\sin\phi$: it's always positive.