Math 32BH: Calculus of Several Variables II (Honors)

Lecture Notes
Austin Christian

Winter 2021

Abstract

These are the class notes for Math 32BH at UCLA in Winter 2021. These notes are meant to supplement our class meetings and textbook readings. In particular, just reading the notes is a very poor substitute for attending class and office hours.

There's a very good chance that you can find typos or more serious errors in these notes — if so, please let me know.
Last update: March 20, 2021

Contents

Acknowledgments. Thanks to the students of 32BH, as well as the TA Frederick Vu, for a fun quarter, and for catching lots of typos (and more than a few serious errors)! It should be noted that each of the following texts influenced these notes to some degree:

- *A Geometric Approach to Differential Forms*, Bachman
- *Advanced Calculus*, Taylor
- *Calculus*, Anton-Bivens-Davis
- *Calculus an Analysis in Euclidean Space*, Shurman
- *Differential Topology*, Guillemin-Pollack
- *Elementary Differential Geometry*, O'Neill
- *Geometry of Differential Forms*, Morita
- *Multivariable Calculus*, Rogawski-Adams-Franzosa

1 Week 1

Our goal this week is to define the (Riemann) integral of a scalar-valued function $f : \mathbb{R}^n \to \mathbb{R}$, and to develop the ability to discern between functions which are integrable and those which are not. This week will be — by a long shot — the most tedious of the quarter, and we'll really get into the weeds. If you find that you enjoy the careful arguments we take up this week, then you can look forward to analysis courses such as 131ABC and 132H; if ϵ really isn't your letter, just hang on until next week, when we take a more computational approach to the material.

1.1 Day 1: Preliminaries

Goals

By the end of today's class, we should be able to do the following.

1. Identify **compact boxes** in $\mathbb{R}^n$, as well as **partitions** of compact boxes.
2. Determine the **infimum** and **supremum** of a set of real numbers, and say what it means for a function to be **bounded**.
3. Interpret the **lower sum** and **upper sum** of a bounded function over a partition as an area or volume.

1.1.1 Some vocabulary in $\mathbb{R}$

Definition. A **compact interval** is a subset of $\mathbb{R}$ of the form

$$[a, b] := \{x \in \mathbb{R} : a \leq x \leq b\},$$

where a and b are real numbers. A **compact subinterval** of $[a, b]$ is a compact interval $[c, d]$ which is contained in $[a, b]$.

Remark. Notice that the empty set is a compact interval (we could take $a > b$), as are singleton sets $\{a\}$ (corresponding to $a = b$).

Definition. The **length** of a compact interval I, denoted $|I|$, is defined as follows. If $I = [a, b]$ for some real numbers $a < b$, then $|I| := b - a$. Otherwise — if I is a singleton or the empty set — $|I| := 0$.

Definition. Let $I = [a, b]$ be a compact interval. A **partition** of I is a finite set of real numbers $\mathcal{P} = \{x_0, x_1, \ldots, x_k\}$ such that

$$a = x_0 < x_1 < \cdots < x_k = b.$$

The **subintervals corresponding to the partition** $\mathcal{P}$ are given by

$$J_i = [x_{i-1}, x_i] \subset [a, b] = I,$$

for $i = 1, \ldots, k$.

Remark.

1. Notice that the empty set does not admit any partitions, and a singleton set admits only one partition. So the only intervals with an interesting set of partitions are those with nonzero length.
2. Because a partition is a finite set, each subinterval associated to a partition has nonzero length.

Exercise 1.1. Suppose $\mathcal{P}$ is a partition of the compact interval I, with corresponding subintervals $J_1, \ldots, J_k$. Prove that

$$|I| = \sum_{i=1}^{k} |J_i|.$$

Definition. Let A be a subset of $\mathbb{R}$, and let $f : A \to \mathbb{R}$ be a function. We say that f is **bounded** if its image $\{f(x) : x \in A\}$ is a subset of some compact interval.

Example 1.2. The functions $f(x) = e^{-x^2}$ and $g(x) = \arctan(x)$ are bounded; notice that the domain for each of these functions is $\mathbb{R}$ — which is *not* bounded. On the other hand, the function $\tan\colon (-\pi/2, \pi/2) \to \mathbb{R}$ is defined on a bounded interval, but is not bounded.

Exercise 1.3. Prove that if I is a compact interval and $f : I \to \mathbb{R}$ is continuous, then f is bounded.

Before continuing our march towards a definition of the integral, we must quickly recall the notions of *least upper bound* and *greatest lower bound* for a subset of the real numbers.

Definition. Let A be a subset of $\mathbb{R}$, and let M be a real number. We say that M is an **upper bound** for A if we have $x \leq M$ for every $x \in A$. We say that M is a **least upper bound** for A if we additionally know that any other upper bound N for A satisfies $N \geq M$.

The following result, whose proof we omit, establishes the existence and uniqueness of a least upper bound.

Theorem 1.4

Let A be a nonempty subset of $\mathbb{R}$. If A has an upper bound, then A has exactly one least upper bound.

This theorem justifies the following definition.

Definition. Let A be a subset of the real numbers. If A is non-empty and has an upper bound, define $\sup(A)$ to be the least upper bound of A. If A is non-empty and has no upper bound, define $\sup(A) := +\infty$. Finally, if A is empty, define $\sup(A) := -\infty$. We call $\sup(A)$ the **supremum** of A.

Analogously, we may define the notions of *greatest lower bound* and *infimum*.

Exercise 1.5. Produce definitions for the terms *lower bound* and *greatest lower bound*. State a result for greatest lower bounds which is analogous to Theorem 1.4. Finally, give a definition for the *infimum*.

Remark. Try not to be intimidated by this new language. To guide your intuition, it's relatively harmless to think of *sup* as *max* and *inf* as *min*. The distinction is that a set is not required to contain its supremum or infimum.

You'll recall that the *definite integral* of a function $f : [a, b] \to \mathbb{R}$ over $[a, b]$ is meant to measure the area of the region bounded by $x = a$, $x = b$, $y = 0$, and $y = f(x)$ in the xy-plane. We can now give our first approximations of this area.

Definition. Let I be a nonempty compact interval, and let $f : I \to \mathbb{R}$ be a bounded function. For every nonemtpy compact subinterval J of I, introduce the notation

$$m_J(f) := \inf\{f(x) : x \in J\} \quad \text{and} \quad M_J(f) := \sup\{f(x) : x \in J\}.$$

Now let $\mathcal{P}$ be a partition of I, with corresponding subintervals $J_1, \ldots, J_k$. The **lower sum** of f over $\mathcal{P}$ is

$$L(f, \mathcal{P}) := \sum_{i=1}^{k} m_{J_i}(f)|J_i|,$$

and the **upper sum** of f over $\mathcal{P}$ is

$$U(f, \mathcal{P}) := \sum_{i=1}^{k} M_{J_i}(f)|J_i|.$$

Example 1.6. Consider the function $f(x) = 1 - x^2$ on the interval $I = [-1, 1]$. Let's compute the lower and upper sums of f over the partition $\mathcal{P} = \{-1, -1/2, 0, 1/3, 1\}$. The corresponding subintervals are

$$J_1 = [-1, -1/2], \quad J_2 = [-1/2, 0], \quad J_3 = [0, 1/3], \quad J_4 = [1/3, 1],$$

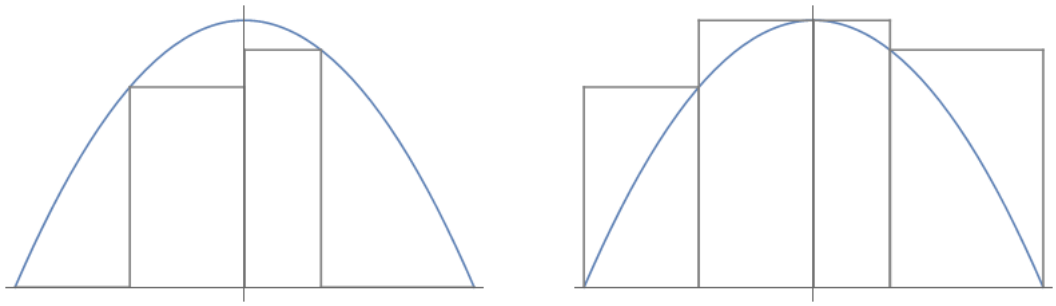

(a) $L(f,\mathcal{P})$ (b) $U(f,\mathcal{P})$

Figure 1.1: Graphical interpretations of the lower and upper sums for the function $f(x) = 1 - x^2$ over the partition $\mathcal{P} = \{-1, -1/2, 0, 1/3, 1\}$.

with lengths given by

$$|J_1| = \frac{1}{2}, \quad |J_2| = \frac{1}{2}, \quad |J_3| = \frac{1}{3}, \quad |J_4| = \frac{2}{3}.$$

It's then not too difficult to compute our extreme values:

$$m_{J_1}(f) = 0, \quad m_{J_2}(f) = \frac{3}{4}, \quad m_{J_3}(f) = \frac{8}{9}, \quad m_{J_4}(f) = 0,$$

while

$$M_{J_1}(f) = \frac{3}{4}, \quad M_{J_2}(f) = 1, \quad M_{J_3}(f) = 1, \quad M_{J_4}(f) = \frac{8}{9}.$$

Finally, we have

$$L(f,\mathcal{P}) = \sum_{i=1}^{4} m_{J_i}(f)|J_i| = (0)\left(\frac{1}{2}\right) + \left(\frac{3}{4}\right)\left(\frac{1}{2}\right) + \left(\frac{8}{9}\right)\left(\frac{1}{3}\right) + (0)\left(\frac{2}{3}\right) = \frac{145}{216}$$

for our lower sum and

$$U(f,\mathcal{P}) = \sum_{i=1}^{4} M_{J_i}(f)|J_i| = \left(\frac{3}{4}\right)\left(\frac{1}{2}\right) + (1)\left(\frac{1}{2}\right) + (1)\left(\frac{1}{3}\right) + \left(\frac{8}{9}\right)\left(\frac{2}{3}\right) = \frac{389}{216}$$

for our upper sum. These sums are depicted in Figure 1.1.

1.1.2 The same vocabulary in $\mathbb{R}^n$

In this subsection, we'll repeat many of the definitions seen above, replacing $\mathbb{R}$ with $\mathbb{R}^n$. For instance, an interval in $\mathbb{R}$ is analogous to a box in $\mathbb{R}^n$.

Definition. A **compact box** in $\mathbb{R}^n$ is a Cartesian product

$$B = I_1 \times I_2 \times \cdots \times I_n$$

of compact intervals $I_1, \ldots, I_n$. The **volume** of the box is the product of the lengths of its sides:

$$|B| := \mathrm{vol}(B) := \prod_{j=1}^{n} |I_j|.$$

A **compact subbox** of B is a compact box which is contained in B.

Remark. Notice that, according to this definition, a compact box must have its sides parallel to the coordinate axes of $\mathbb{R}^n$. So, for instance, the compact square $\{(x,y) \in \mathbb{R}^2 : |x| + |y| \leq 1\}$ in $\mathbb{R}^2$ does not satisfy this definition of compact box.

Just as we partition compact intervals of $\mathbb{R}$, we may partition compact boxes in $\mathbb{R}^n$. We do this by considering a product of partitions.

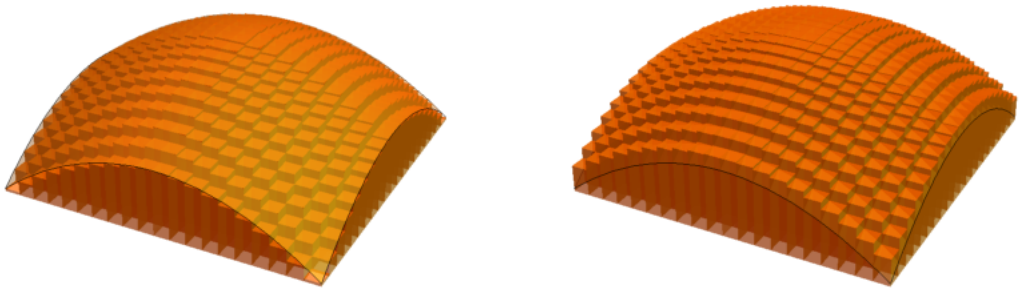

(a) $L(f,\mathcal{P})$ (b) $U(f,\mathcal{P})$

Figure 1.2: Graphical interpretations of the lower and upper sums for a function of two variables, with a fine partition.

Definition. A **partition** of a box $B = I_1 \times I_2 \times \cdots \times I_n$ in $\mathbb{R}^n$ is a Cartesian product

$$\mathcal{P} = \mathcal{P}_1 \times \mathcal{P}_2 \times \cdots \times \mathcal{P}_n$$

of partitions $\mathcal{P}_j$ for I_j.

Notice that there are many subboxes J of B associated to a partition $\mathcal{P}$. In particular, if $J_{i_1}, \ldots, J_{i_n}$ are subintervals of $I_1, \ldots, I_n$ corresponding to $\mathcal{P}_1, \ldots, \mathcal{P}_n$, respectively, then

$$J_{i_1} \times \cdots \times J_{i_n}$$

is a subbox of B associated to $\mathcal{P}$. Coming up with notation for these subboxes doesn't sound very fun (we have n partitions to think about, and these partitions need not have the same cardinality), so we will just refer to these as **subboxes of** $\mathcal{P}$.

Remark. Just as a subinterval associated to a partition of a compact interval could not have length zero, a subbox associated to a partition of a compact box cannot have volume zero.

After partitions, we defined the notion of boundedness for functions of a single variable. The notion in higher dimensions is precisely the same.

Definition. Let A be a subset of $\mathbb{R}^n$, and let $f : A \to \mathbb{R}$ be a function. We say that f is **bounded** if its image $\{f(x) : x \in A\}$ is a subset of some compact interval.

At last, we can define the lower and upper sums in any dimension.

Definition. Let B be a nonempty compact box in $\mathbb{R}^n$, and let $f : B \to \mathbb{R}$ be a bounded function. For every nonemtpy compact subbox J of B, introduce the notation

$$m_J(f) := \inf\{f(x) : x \in J\} \quad \text{and} \quad M_J(f) := \sup\{f(x) : x \in J\}.$$

Now let $\mathcal{P}$ be a partition of B, with associated subboxes J. The **lower sum** and **upper sum** of f over $\mathcal{P}$ are

$$L(f,\mathcal{P}) := \sum_J m_J(f)|J| \quad \text{and} \quad U(f,\mathcal{P}) := \sum_J M_J(f)|J|.$$

respectively, where the sum is taken over all boxes J associated to $\mathcal{P}$.

Just as in the case $n = 1$, the lower and upper sums admit a natural interpretation graphically. An example in the case $n = 2$ can be seen in Figure 1.2. Notice that the lower sum uses boxes which live entirely below the (translucent) graph of f, while the upper sum uses boxes which lie entirely above the graph.

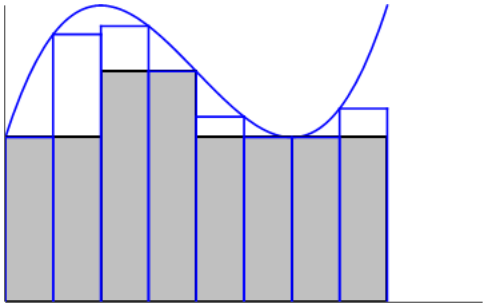

Figure 1.3: A finer partition leads to a larger lower sum.

1.2 Day 2: The Riemann integral over boxes

Goals

By the end of today's class, we should be able to do the following.

1. State the definition of the **lower integral** and **upper integral** of a bounded function over a compact interval.
2. Give examples of **integrable** and **nonintegrable** functions.
3. Use the **integrability criterion** to determine whether or not a function is integrable.

Lemma 1.7. *For every compact box B, every partition $\mathcal{P}$ of B, and every bounded function $f : B \to \mathbb{R}$, $L(f,\mathcal{P}) \leq U(f,\mathcal{P})$.*

Exercise 1.8. Prove Lemma 1.7

Definition. Let $\mathcal{P}$ and $\mathcal{P}'$ be partitions of a compact box B. Then $\mathcal{P}'$ is a **refinement** of $\mathcal{P}$ if $\mathcal{P} \subset \mathcal{P}'$, in which case we say that $\mathcal{P}'$ **refines** $\mathcal{P}$.

Lemma 1.9. *Suppose that $\mathcal{P}$ and $\mathcal{P}'$ are partitions of B, and that $\mathcal{P}'$ is a refinement of $\mathcal{P}$. Then*

$$L(f,\mathcal{P}) \leq L(f,\mathcal{P}') \quad \textit{and} \quad U(f,\mathcal{P}') \leq U(f,\mathcal{P}).$$

Idea of proof. We'll give a proof-by-picture for the first claim in the case $n = 1$. The basic idea is that when we take a finer partition, our boxes are smaller, and thus $m_J(f)$ has less of an opportunity to be small, so the lower sum is larger. Similarly, $M_J(f)$ has less opportunity to be large, so the upper sum is smaller. Consider Figure 1.3. The gray boxes depict the lower sum for f over a partition $\mathcal{P}$; we may refine this partition by cutting each interval in half. This produces the refinement $\mathcal{P}'$ of $\mathcal{P}$, and we see that this partition leads to a larger lower sum. □

Exercise 1.10. Turn the proof-by-picture into a formal proof.

Definition. Given two partitions

$$\mathcal{P} = \mathcal{P}_1 \times \cdots \times \mathcal{P}_n \quad \text{and} \quad \mathcal{P}' = \mathcal{P}'_1 \times \cdots \times \mathcal{P}'_n$$

of a compact box B in $\mathbb{R}^n$, the **common refinement** of $\mathcal{P}$ and $\mathcal{P}'$ is the partition

$$\mathcal{P}'' = (\mathcal{P}_1 \cup \mathcal{P}'_1) \times \cdots \times (\mathcal{P}_n \cup \mathcal{P}'_n).$$

The point of the common refinement is that it is the smallest partition which refines both $\mathcal{P}$ and $\mathcal{P}'$. Intuitively, a partition tells us where to cut a box into subboxes; the common refinement tells us to cut wherever either of $\mathcal{P}$ or $\mathcal{P}'$ cuts.

We are now ready to say what it means for a bounded function on a nonempty compact box in $\mathbb{R}^n$ to be integrable.

Definition. Let $f : B \to \mathbb{R}$ be a bounded function, where B is a nonempty compact box in $\mathbb{R}^n$. The **lower integral** of f over B, denoted $L\int_B f$, is the least upper bound of the lower sums of f over all partitions of B:

$$L\int_B f := \sup\{L(f,\mathcal{P}) : \mathcal{P} \text{ is a partition of } B\}.$$

Similarly, the **upper integral** of f over B is denoted $U\int_B f$, and is the greatest lower bound of the upper sums of f over all partitions of B:

$$U\int_B f := \inf\{U(f,\mathcal{P}) : \mathcal{P} \text{ is a partition of } B\}.$$

We say that f is **integrable** over B if $L\int_B f = U\int_B f$, and in this case we define the **integral** of f over B to be their common value: $\int_B f := L\int_B f = U\int_B f$.

Yikes. This definition is sort of a lot to take in, but our graphical interpretations of upper and lower sums can help us develop an intuition. Lower sums always give an underestimate of the value we're after, and the lower integral is the *biggest such underestimate*[1]. In the $n = 1$ case, this means that the area under our curve is at least as large as $L\int_B f$. Upper sums, on the other hand, overestimate the value we're trying to compute, and the upper integral is the *smallest such overestimate*[2]. So the true value of the integral is no more than $U\int_B f$. If we have

$$L\int_B f = U\int_B f,$$

then the value of $\int_B f$ must be exactly this common value. If, on the other hand, $L\int_B f < U\int_B f$, then we can't pinpoint a precise value for $\int_B f$, and we say that f is not integrable over B.

Lemma 1.11. *Let B be a compact box in $\mathbb{R}^n$, $f : B \to \mathbb{R}$ a bounded function. Then the lower integral of f over B is no larger than the upper integral of f over B.*

Exercise 1.12. Prove Lemma 1.11.
Hint: First show that if L and U are subsets of $\mathbb{R}$ with the property that $\ell \le u$ for all $\ell \in L$ and $u \in U$, then $\sup(L) \le \inf(U)$.

Corollary 1.13. *Let B be a compact box in $\mathbb{R}^n$, $f : B \to \mathbb{R}$ a bounded function. Then f is integrable over B if and only if $L\int_B f \ge U\int_B f$.*

Proof. Lemma 1.11 tells us that $L\int_B f \le U\int_B f$, so $L\int_B f = U\int_B f$ if and only if $L\int_B f \ge U\int_B f$. □

Except in particularly boring cases — constant functions, step functions, linear functions, etc. — computing an integral from the definition is a far-from-straightforward task. You may recall computing some Riemann sums in your single-variable calculus class before learning about the Fundamental Theorem of Integral Calculus, and you probably had a bad time. Our story will follow a similar arc this quarter, and we'll soon meet a savior that spares us the pain of computing integrals from the definition.

First, however, let's meet a character you probably missed in your first encounter with the integral: a non-integrable function.

Example 1.14. In this example, we will define a function using the *rational numbers* $\mathbb{Q}$. Recall that $\mathbb{Q}$ consists of those real numbers which can be expressed as a ratio of the form p/q, where p and q are integers, and q is nonzero. The relevant fact about $\mathbb{Q}$ for this example is that every interval in $\mathbb{R}$ contains rational numbers as well as irrational numbers — those real numbers which are not in $\mathbb{Q}$. For this reason we say that both $\mathbb{Q}$ and its complement $\mathbb{R} \setminus \mathbb{Q}$ are *dense* in $\mathbb{R}$.

[1] This is a little bit of a lie. It may be the case that none of the lower sums are actually equal to the lower integral, just that we can find lower sums arbitrarily close to the lower integral. This is the difference between max and sup, but this isn't a crucial distinction when developing intuition.

[2] This is also a lie, since min ≠ inf.

Now let $B = [0,1]$, and define $f : B \to \mathbb{R}$ by

$$f(x) := \begin{cases} 1, & x \in \mathbb{Q} \\ 0, & x \notin \mathbb{Q} \end{cases}.$$

This is known as the *Dirichlet function*, or the *indicator function of* $\mathbb{Q}$. I've also heard it called *the salt and pepper function*, but I don't see that catching on. Now suppose that J is a compact subinterval of B of nonzero length. Then J contains points which are in $\mathbb{Q}$ as well as points which are not in $\mathbb{Q}$. For this reason,

$$m_J(f) = 0 \quad \text{and} \quad M_J(f) = 1.$$

Then for any partition $\mathcal{P}$ of B we have

$$L(f,\mathcal{P}) = \sum_J m_J(f)|J| \quad \text{and} \quad U(f,\mathcal{P}) = \sum_J M_J(f)|J|,$$

where the sums are taken over all compact subintervals J corresponding to $\mathcal{P}$. Based on our computation of $m_J(f)$ and $M_J(f)$, these turn into[3]

$$L(f,\mathcal{P}) = \sum_J 0 \cdot |J| = 0 \quad \text{and} \quad U(f,\mathcal{P}) = \sum_J 1 \cdot |J| = |B| = 1.$$

(The fact that the lengths of the subintervals sum to the length of B follows from Exercise 1.1.) That is, every lower sum of f is 0, and every upper sum of f is 1. So

$$L\int_B f = 0 \quad \text{and} \quad U\int_B f = 1.$$

Since $L\int_B f \neq U\int_B f$, we see that f is not integrable over B.

Remark. It's more accurate to say that the Dirichlet function is not *Riemann* integrable. A little later in your mathematical career you'll learn about the *Lebesgue integral*, which is defined for this function, and evaluates to 0.

Exercise 1.15. Letting f denote the Dirichlet function from the previous example, construct a sequence $\{f_n\}$ of functions $f_n : [0,1] \to \mathbb{R}$, each of which is integrable over $[0,1]$, with the property that $\lim_{n\to\infty} f_n(x) = f(x)$ for every $x \in [0,1]$. This demonstrates the fact that Riemann integrability is not preserved under pointwise convergence.
Hint: Because $\mathbb{Q}$ is countable, we may choose an enumeration $\{q_n\}$ of $\mathbb{Q}$. Now "turn on" one rational number at a time.

We've already said that we won't be doing much computing with the definition of the integral, but even checking whether or not a function is integrable can be an annoying task. For instance, a polynomial function seems like it *should* be integrable, but what's your proof of this fact? The following result will be helpful over and over again in showing that a function is integrable.

Theorem 1.16: The integrability criterion

Let B be a compact box in $\mathbb{R}^n$, $f : B \to \mathbb{R}$ a bounded function. Then f is integrable over B if and only if, for every $\epsilon > 0$, there exists a partition $\mathcal{P}$ of B such that $U(f,\mathcal{P}) - L(f,\mathcal{P}) < \epsilon$.

Proof. We will prove the portion of this theorem that we will use most frequently. Namely, we'll show that if, for every $\epsilon > 0$, there exists a partition $\mathcal{P}$ of B satisfying $U(f,\mathcal{P}) - L(f,\mathcal{P}) < \epsilon$, then f is integrable. To show that f is integrable, we need to show that $U\int_B f = L\int_B f$ or, equivalently, that

$$U\int_B f - L\int_B f = 0.$$

We'll do this by showing that $U\int_B f - L\int_B f$ is smaller than any positive number. Take some arbitrarily small $\epsilon > 0$. By assumption, we can find $\mathcal{P}$ satisfying

$$U(f,\mathcal{P}) - L(f,\mathcal{P}) < \epsilon.$$

[3]You might contest that our computation of $m_J(f)$ and $M_J(f)$ does not apply to intervals with zero length. But recall that a subinterval associated to a partition cannot have zero length.

Now $U(f,\mathcal{P}) \geq U\int_B f$ and $L(f,\mathcal{P}) \leq L\int_B f$, so this implies that

$$U\int_B f - L\int_B f < \epsilon.$$

So $U\int_B f - L\int_B f$ is smaller than any positive number. Lemma 1.11 tells us that $U\int_B f - L\int_B f$ cannot be negative, so this quantity must be zero, as desired. □

Because it only requires a single partition $\mathcal{P}$ for a given real number $\epsilon > 0$, Theorem 1.16 allows us to avoid thinking about the infimum and supremum involved in the definition of the upper and lower integral.

Example 1.17. Let's use the integrability criterion to show that the function $f : [0,1] \to \mathbb{R}$ defined by

$$f(x) := \begin{cases} 0, & 0 \leq x \leq 1/2 \\ 8, & 1/2 < x \leq 1 \end{cases}$$

is integrable. To do this, we start with an arbitrary $\epsilon > 0$, and we need to construct a partition $\mathcal{P}$ of $[0,1]$ with the property that $U(f,\mathcal{P}) - L(f,\mathcal{P}) < \epsilon$. We can do this by squeezing the "jump" which occurs at $x = 1/2$ into a very small interval. Concretely, suppose $h > 0$ is some number smaller than $1/2$. Then we can consider the partition

$$\mathcal{P}_h = \{0, 1/2-h, 1/2+h, 1\}.$$

The subintervals associated to this partition are

$$J_1 = [0, 1/2-h], \quad J_2 = [1/2-h, 1/2+h], \quad \text{and} \quad J_3 = [1/2+h, 1].$$

The extreme values on these intervals are given by

$$m_{J_1}(f) = 0, \quad m_{J_2}(f) = 0, \quad \text{and} \quad m_{J_3}(f) = 8,$$

and

$$M_{J_1}(f) = 0, \quad M_{J_2}(f) = 8, \quad \text{and} \quad M_{J_3}(f) = 8.$$

So the lower sum is

$$L(f,\mathcal{P}_h) = (0)(1/2-h) + (0)(2h) + (8)(1/2-h) = 4 - 8h$$

and the upper sum is

$$U(f,\mathcal{P}_h) = (0)(1/2-h) + (8)(2h) + (8)(1/2-h) = 4 + 8h.$$

The difference between these is

$$U(f,\mathcal{P}_h) - L(f,\mathcal{P}_h) = 16h.$$

The integrability criterion tells us that f is integrable if we can produce a partition so that this difference is less than ϵ, for any given $\epsilon > 0$. Given $\epsilon > 0$, we can simply choose h smaller than $\epsilon/16$, and our work here shows that $U(f,\mathcal{P}_h) - L(f,\mathcal{P}_h) < \epsilon$. So f is indeed integrable.

1.3 Day 3: Continuity and integrability

Goals

By the end of today's class, we should be able to do the following.

1. Say what it means for a function of several variables to be **continuous** — both via the formal definition and via informal intuition.
2. Explain (at least informally) why continuous functions on compact boxes are integrable.

Definition. Let A be a subset of $\mathbb{R}^n$, and consider a function $f : A \to \mathbb{R}$. We say that f is **continuous at** $\vec{x}_0 \in A$ if, for every $\epsilon > 0$, we may choose $\delta > 0$ such that if $\vec{x}_1 \in A$ is an element of A satisfying $d(\vec{x}_0, \vec{x}_1) < \delta$, then we have $|f(\vec{x}_0) - f(\vec{x}_1)| < \epsilon$. If f is continuous at $\vec{x}$ for every $\vec{x} \in A$, we simply say that f is **continuous**.

The idea you should internalize is this: for continuous functions, a small change in the input leads to a small change in the output.

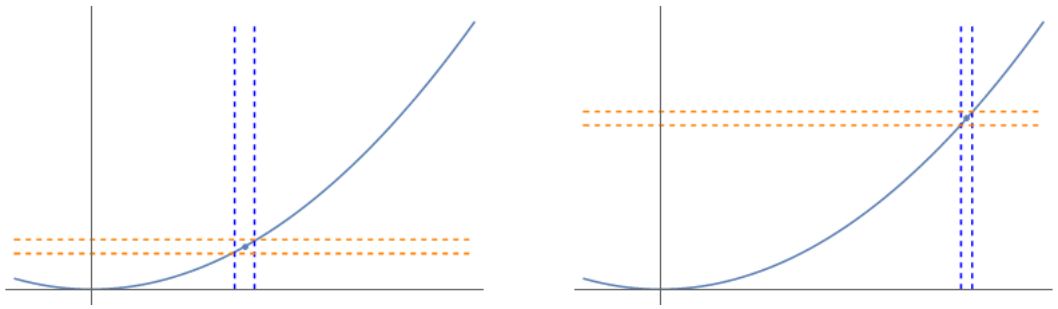

(a) Smaller x_0 (b) Larger x_0

Figure 1.4: For the function $f(x) = x^2$, different values of x_0 require different values of δ — even for the same ϵ.

Example 1.18. Consider the function $f : \mathbb{R}^2 \to \mathbb{R}$ defined by

$$f(x,y) := \begin{cases} 1, & x = 0 \\ -1, & x \neq 0 \end{cases}.$$

This function certainly doesn't match our intuitive notion of continuity — there are inputs very close to the line $x = 0$ whose output is -1, while inputs on this line have output 1. To prove that f is not continuous, we simply need to exhibit some $\epsilon > 0$ (in this case, we get to choose $\epsilon > 0$) such that, no matter how small $\delta > 0$ is chosen (we don't get to choose $\delta > 0$), there are points (x_1, y_1) and (x_2, y_2) whose distance from each other is less than δ, but whose function values differ by more than ϵ.

Concretely, suppose $\epsilon = 1/2$, and let $\delta > 0$ be arbitrary. Consider the point $(\delta/2, 0)$. The distance between this point and $(0,0)$ is less than δ:

$$d((\delta/2, 0), (0,0)) = \delta/2 < \delta.$$

At the same time,

$$|f(\delta/2, 0) - f(0,0)| = |-1-1| = 2 > \frac{1}{2} = \epsilon.$$

So we can choose points which are arbitrarily close together whose function values are not close to one another, meaning that f is not continuous at $(0,0)$. (Indeed, f fails to be continuous at any point with $x = 0$.)

You may remember (depending on where you learned calculus) having to use the definition of continuity to prove that certain functions are continuous. No such problems will be formally assigned in this course, but we will discuss one example of this sort in order to better understand **uniform continuity**.

Example 1.19. Consider the function $f : \mathbb{R} \to \mathbb{R}$ defined by $f(x) = x^2$. This satisfies our intuitive notion of continuity, so let's think about how we would check the definition. Namely, we want to show that f is continuous at x_0, where x_0 is an arbitrary real number. Say we're given[4] $\epsilon > 0$. We need to come up with $\delta > 0$ so that if $|x - x_0| < \delta$, then $|f(x) - f(x_0)| < \epsilon$. Actually doing this isn't hard, but it's kind of tedious.

Instead, let's think about the graph of f, shown in Figure 1.4. In each of Figures 1.4a and 1.4b we have selected a value for x_0; the chosen value for $\epsilon > 0$ is the same in both figures, depicted via the orange lines centered on $y = f(x_0)$. Our goal is to choose a $\delta > 0$ such that if $|x - x_0| < \delta$, then $|f(x) - f(x_0)| < \epsilon$. We can find such a δ for each value of x_0 shown, as shown by the blue lines centered on $x = x_0$. But notice that the δ we choose for the larger value of x_0 is necessarily smaller. In words, we have some threshold for acceptable error in our output (measured by ϵ), and to keep our output error under this threshold, we can only allow so much error in our input (measured by δ). Because we can produce a $\delta > 0$ for any choice of $\epsilon > 0$ and $x_0 \in \mathbb{R}$, f is continuous on all of $\mathbb{R}$.

We must point out, however, that the amount of error we can allow in our input in the previous example depends on x_0 — and as x_0 grows without bound, the acceptable input error shrinks to 0. In particular, given $\epsilon > 0$, we cannot choose a single $\delta > 0$ that works for all values of x_0. This is okay — our function is still continuous — but it leads us to the following, greedier notion of continuity.

[4] When proving that a function is continuous, we're given $\epsilon > 0$ and we choose $\delta > 0$, the opposite of the situation when we're proving that a function is not continuous.

Definition. Let A be a subset of $\mathbb{R}^n$, and consider a function $f : A \to \mathbb{R}$. We say that f is **uniformly continuous** if, for every $\epsilon > 0$, we may choose $\delta > 0$ such that if $\vec{x}_1, \vec{x}_2 \in A$ are an elements of A satisfying $d(\vec{x}_1, \vec{x}_2) < \delta$, then we have $|f(\vec{x}_1) - f(\vec{x}_2)| < \epsilon$.

Remark. Notice the distinction between uniform continuity and continuity: the definition of continuity allows us to choose $\delta > 0$ based on both $\epsilon > 0$ and a point $\vec{x}_0 \in A$, while uniform continuity demands a δ which depends only on ϵ, and works for any choice of $\vec{x}_0 \in A$. Clearly any function which is uniformly continuous is continuous.

Exercise 1.20. Let $f : (0, \infty) \to \mathbb{R}$ be defined by $f(x) := 1/x$. Is f uniformly continuous? What if we restrict f to $(0, 1]$? What about $[1, \infty)$?

One of our goals for today is to show that continuous functions on compact boxes are integrable. We restrict attention to functions on compact boxes in order to take advantage of the following fact[5].

Lemma 1.21. *Let B be a compact box in $\mathbb{R}^n$, and let $f : B \to \mathbb{R}$ be a continuous function. Then f is uniformly continuous.*

Exercise 1.22. Read a proof of Lemma 1.21 (or write your own!).
Note: The proof of this fact gets at the heart of the notion of compactness, so you may like thinking about it if analysis is your thing. But it's certainly not something you're responsible for in this course.

Taking Lemma 1.21 as true, today's primary goal is well within reach.

Theorem 1.23: Continuity implies integrability

Let B be a compact box in $\mathbb{R}^n$. If $f : B \to \mathbb{R}$ is continuous, then f is integrable.

Proof. We'll prove this theorem using Theorem 1.16, the integrability criterion. So we need to show that, given some arbitrary $\epsilon > 0$, we can find a partition $\mathcal{P}$ of B with the property that

$$U(f, \mathcal{P}) - L(f, \mathcal{P}) < \epsilon.$$

First, let's make an obvious observation: if $\text{vol}(B) = 0$, then f is automatically integrable — all of the lower sums and upper sums are just 0, so the lower and upper integrals agree. With this out of the way, we can assume that $\text{vol}(B) \neq 0$ and define

$$\epsilon' := \frac{\epsilon}{\text{vol}(B)}.$$

Next, Lemma 1.21 tells us that f is uniformly continuous. This means that we can find $\delta > 0$ with the property that whenever $\vec{x}_1, \vec{x}_2 \in B$ have $d(\vec{x}_1, \vec{x}_2) < \delta$, we have $|f(\vec{x}_1) - f(\vec{x}_2)| < \epsilon'$.

Now comes the first step that requires a dose of geometric reasoning. Let's take a partition $\mathcal{P}$ of B so that every subbox J associated to $\mathcal{P}$ has side lengths less than δ/n. This ensures that if $\vec{x}_1, \vec{x}_2 \in J$ for some subbox J associated to $\mathcal{P}$, then $d(\vec{x}_1, \vec{x}_2) < \delta$ (why?). But then, by our choice of δ, we know that

$$|f(\vec{x}_1) - f(\vec{x}_2)| < \epsilon' = \frac{\epsilon}{\text{vol}(B)}.$$

In particular, the supremum of all values attained by f as $\vec{x}$ varies over J can differ from the infimum by at most ϵ':

$$M_J(f) - m_J(f) < \epsilon' = \frac{\epsilon}{\text{vol}(B)}.$$

Finally, we multiply by $\text{vol}(J)$ and sum over all subboxes J corresponding to $\mathcal{P}$:

$$\left(\sum_J M_J(f)\,\text{vol}(J)\right) - \left(\sum_J m_J(f)\,\text{vol}(J)\right) = \sum_J \frac{\epsilon\,\text{vol}(J)}{\text{vol}(B)} = \frac{\epsilon}{\text{vol}(B)} \sum_J \text{vol}(J).$$

The two sums on the left hand side of this inequality are the upper and lower sums of f over $\mathcal{P}$, respectively. Also, $\sum_J \text{vol}(J) = \text{vol}(B)$, so we have

$$U(f, \mathcal{P}) - L(f, \mathcal{P}) < \epsilon.$$

According to the integrability criterion, we may conclude that f is integrable. □

[5] We haven't said what it means for a subset of $\mathbb{R}^n$ to be compact, but the fact actually holds for any compact subset of $\mathbb{R}^n$ — not just boxes.

Exercise 1.24. Let J be a compact box in $\mathbb{R}^n$ whose sides each have length less than δ/n. Show that if $\vec{x}_1$ and $\vec{x}_2$ are a pair of points in J, then $d(\vec{x}_1, \vec{x}_2) < \delta$.

2 Week 2

We begin this week by finishing up our discussion of the fact that continuity implies integrability (see notes above). We will then spend some time defining integration over subsets of $\mathbb{R}^n$ which are not compact boxes, but with far less attention to detail than we employed last week. We then turn to the star of this week: Fubini's theorem. This theorem will allow us to begin evaluating integrals of functions of several variables using familiar techniques from single-variable calculus. After sketching the proof of Fubini's theorem, we will practice using the result by seeing several examples.

2.1 Day 4: Integration over nonboxes

Goals

By the end of today's class, we should be able to do the following.

1. Give an informal explanation of how to define the integral of a function $f : K \to \mathbb{R}$, where K is a compact subset of $\mathbb{R}^n$ whose boundary has volume zero.
2. State the difference between a **multiple integral** and an **iterated integral**.

2.1.1 Defining integration over nonboxes

Definition. Let S be a subset of $\mathbb{R}^n$. The **indicator function** of S is a function $\chi_S : \mathbb{R}^n \to \mathbb{R}$ defined by

$$\chi_S(x) := \begin{cases} 1, & x \in S \\ 0, & x \notin S \end{cases}.$$

The indicator function is also called the **characteristic function** of S.

For many bounded subsets of $\mathbb{R}^n$, we can use the indicator function to define the set's volume.

Definition. The **volume** of a bounded subset S of $\mathbb{R}^n$ is defined to be

$$\mathrm{vol}(S) := \int_B \chi_S,$$

where B is any compact box in $\mathbb{R}^n$ containing S, *if this integral exists*.

Exercise 2.1. Prove that volume is well-defined. i.e., show that if $\int_B \chi_S$ exists for some compact box B containing S and B' is another compact box containing S, then $\int_{B'} \chi_S$ exists and $\int_{B'} \chi_S = \int_B \chi_S$.

Exercise 2.2. Find a nonempty subset of $\mathbb{R}$ whose volume is not defined.

Say we have a compact box B over which we want to integrate some function $f : B \to \mathbb{R}$, and a subset $S \subset B$ which has volume zero. Then, at least heuristically, our integral shouldn't care what f is doing on S. For instance, consider the function $f : [-1, 1] \to \mathbb{R}$ defined by

$$f(x) := \begin{cases} 1, & x \neq 0 \\ 2021, & x = 0 \end{cases}.$$

By any reasonable definition, the integral $\int_{[-1,1]} f$ should be 2 — the fact that f jumps up to 2021 at $x = 0$ shouldn't affect the area under the graph of f, since the singleton $\{x = 0\}$ doesn't have any width. This motivates the definition of **near-continuity**, which asks a function to be continuous, *except perhaps on a subset of volume zero*.

Definition. Let A be a subset of $\mathbb{R}^n$. We call a function $f : A \to \mathbb{R}$ **nearly continuous on** A if there is a subset $S \subset A$ with volume zero so that f is continuous when restricted to $A \setminus S$.

Remark. The basic idea of near-continuity is that "jump discontinuities" are okay from the standpoint of integration. Except, where functions of a single variable "jump" at a single point, a function of two variables

can jump along a point, a line, or a curve, and functions of several variables can jump along any set of volume zero.

Theorem 2.3: Near continuity implies integrability

Let B be a compact box in $\mathbb{R}^n$, and let $f : B \to \mathbb{R}^n$ be a bounded, nearly continuous function. Then f is integrable over B.

Idea of proof. An unfortunate feature of Riemann integration is that the careful proof of Theorem 2.3 is more tedious than the statement initially lets on. Later you'll learn about a better integral where Theorem 2.3 follows more or less immediately from Theorem 1.23. For now, we'll present the basic idea. See the textbook for a careful proof.

The key idea is quite simple: because f is bounded, the only points of B which might present a problem when defining the integral of f are those points where f is discontinuous. But f is *nearly* continuous, so the set of such points has zero volume, and therefore shouldn't create any difference between the lower and upper integrals. We can turn this into a rigorous proof using Theorem 1.16. We are given some $\epsilon > 0$, and we want to find a partition $\mathcal{P}$ satisfying $U(f,\mathcal{P}) - L(f,\mathcal{P}) < \epsilon$. We pull this off by starting with a partition with the following property: let $S \subset B$ be a set of volume zero containing the discontinuities of f. We choose $\mathcal{P}$ so that the sum of the volumes of those subboxes of $\mathcal{P}$ which intersect S is less than $\epsilon/(4R)$, where $R > 0$ is chosen to satisfy $|f(x)| < R$, for all $x \in B$. By covering S with such tiny subboxes, and using the fact that f is bounded, we ensure that any contribution to $L(f,\mathcal{P})$ or $U(f,\mathcal{P})$ made by these subboxes is negligible, and therefore the set of discontinuities doesn't cause f to be non-integrable. □

Definition. Let K be a compact subset of $\mathbb{R}^n$ whose boundary[1] has volume zero, and let $f : K \to \mathbb{R}$ be a bounded, nearly-continuous function. Let B be any compact box in $\mathbb{R}^n$ containing K. Then the **extension of f to B** is the function $\tilde{f} : B \to \mathbb{R}$ defined by

$$\tilde{f}(x) := \begin{cases} f(x), & x \in K \\ 0, & x \notin K \end{cases}.$$

The **integral of f over K** is then defined to be

$$\int_K f := \int_B \tilde{f}.$$

Remark. We haven't defined what it means for a function to be **integrable** over K, because the conditions in the above definition ensure that the integral $\int_B f$ will exist. We're slowly shifting our focus from details such as these towards more computational concerns.

Finally, let's list (without proof) some unsurprising properties of the integral, hopefully familiar from single-variable calculus.

Properties of integrals. Let $K \subset \mathbb{R}^n$ be a compact set whose boundary has volume zero, and let $f, g : K \to \mathbb{R}$ be a pair of nearly continuous functions. Then the following statements are true.

1. (Linearity) For any real numbers $c, d \in \mathbb{R}$, we have $\int_K (cf + dg) = c\int_K f + d\int_K g$.
2. (Monotonicity) If $f \le g$ on K, then $\int_K f \le \int_K g$.
3. (Additivity) Let $K_1, K_2 \subset K$ be a pair of compact subsets of K whose intersection $K_1 \cap K_2$ has volume zero, with $K_1 \cup K_2 = K$. Then $\int_{K_1} f + \int_{K_2} f = \int_K f$.

[1] We won't formally define the *boundary* of a subset of $\mathbb{R}^n$, but for any subset of $\mathbb{R}^n$ we consider in this class, the boundary will be what you think it is, and will have volume zero. The extra words in this definition basically just ensure that we're not trying to integrate over a Cantor set or an Osgood curve, or some other pathological set.

2.1.2 Multiple integrals versus iterated integrals

Now that we've gone through the pain of defining the integrals we care to compute, we'll begin developing some actual integrating skills. The first and most important fact we will learn relates integrals in higher dimensions to iterated versions of the single-variable integrals with which we're familiar. To this end, we introduce some notation. We will use the term **multiple integral** to refer to any integral $\int_R f$, where R is a region in $\mathbb{R}^n$, $n > 1$. In the special cases $n = 2$ and $n = 3$ we say that we have a **double integral** or **triple integral**, respectively.

You may have noticed that our integrals haven't featured any dx or dy terms at their ends. We will change this now, using $\int_R f\,dA$ to denote a double integral and $\int_R f\,dV$ to denote a multiple integral of more variables. Additionally, we will sometimes write

$$\iint_R f\,dA \quad \text{and} \quad \iiint_R f\,dV$$

for the special cases of double and triple integrals. The expressions dA and dV won't have any formal meaning for us right now — though they'll be helpful when performing changes of variables — but hopefully we will be able to revisit these terms later in the quarter.

Finally, we define **iterated integrals**. Iterated integrals are built up from single-variable integrals, and thus can be computed using your existing integration skills. The primary goal of our next class will be to rewrite multiple integrals as iterated integrals.

Say we have a function $f(x, y)$, defined on a box $[a, b] \times [c, d]$ in $\mathbb{R}^2$. Then, for each $x \in [a, b]$, we can think of $f(x, y)$ as a single-variable function on $[c, d]$ by treating x as constant. This allows us to compute

$$\int_c^d f(x, y)dy := \int_{[c,d]} f(x, y),$$

the result of which is a function of x on $[a, b]$. If this function of x is integrable, we can then compute

$$\int_a^b \left(\int_c^d f(x, y)dy \right) dx.$$

This is an iterated integral, because it consists of one one-dimensional integral nested inside of another. More generally, an **n-fold iterated integral** consists of n one-dimensional integrals nested inside one another:

$$\int_{a_1}^{b_1} \int_{a_2(x_1)}^{b_2(x_1)} \cdots \int_{a_n(x_1,\ldots,x_{n-1})}^{b_n(x_1,\ldots,x_{n-1})} f(x_1, x_2, \ldots, x_n)dx_n \cdots dx_2 dx_1.$$

Notice, somewhat confusingly, that the limits of integration of the inner integrals may depend on the variables of the outer variables.

Remark. One reason we've started using terms like dx and dy again is to keep track of the order of our iterated integrals. In the above iterated integral, the term $dx_n \cdots dx_2 dx_1$ tells us that the innermost integral treats x_n as variable, while the variable of the outermost integral is x_1. We could also have an n-fold iterated integral of the form

$$\int_{a_n}^{b_n} \int_{a_{n-1}(x_n)}^{b_{n-1}(x_n)} \cdots \int_{a_1(x_2,\ldots,x_n)}^{b_1(x_2,\ldots,x_n)} f(x_1, x_2, \ldots, x_n)dx_1 \cdots dx_{n-1} dx_n,$$

or have the variables in any other order we like. The important thing is that outer integrals do not depend on inner variables.

Remark. We assume that you have seen (and hopefully proved) the Fundamental Theorem of Integral Calculus before, and will assume this going forward. That is, if $f : [a,b] \to \mathbb{R}$ is a Riemann integrable function on $[a,b]$ and F is an antiderivative of f on $[a,b]$, then

$$\int_a^b f(x)dx := \int_{[a,b]} f = F(b) - F(a).$$

The practical upshot is that we know how to compute integrals of functions of a single variable. We only have this theorem for single-variable integrals, so we'll need to establish a different theorem to compute multiple integrals. This is the topic of our next class meeting.

Example 2.4. Compute the values of the following iterated integrals:

(a) $\int_0^2 \int_0^1 (x^2+y^2)dxdy$

(b) $\int_{-1}^1 \int_0^{\sqrt{1-x^2}} 1 dydx$

Solution.

(a) First we compute the inner integral, $\int_0^1 (x^2+y^2)dx$. The differential dx tells us that we should be treating x as variable and y as constant. We have

$$\int_0^1 (x^2+y^2)dx = \left[\frac{1}{3}x^3 + xy^2\right]_0^1 = \frac{1}{3} + y^2.$$

We can now compute the outer integral:

$$\int_0^2 \left(\frac{1}{3}+y^2\right)dy = \left[\frac{1}{3}y + \frac{1}{3}y^3\right]_0^2 = \frac{2}{3} + \frac{8}{3} = \frac{10}{3}.$$

So $\boxed{\int_0^2 \int_0^1 (x^2+y^2)dxdy = 10/3}$.

(b) Our computation this time will proceed in a manner very similar to (a), except that the upper limit of our inner integral depends on the variable of our outer integral. But this is fine, because we expect the result of our inner integral to be a function of x (in this case) anyway. We have

$$\int_0^{\sqrt{1-x^2}} 1 dy = [y]_0^{\sqrt{1-x^2}} = \sqrt{1-x^2},$$

so the outer integral is $\int_{-1}^1 \sqrt{1-x^2}dx$. Using the substitution $x = \cos\theta$, this becomes

$$\begin{aligned}\int_{-1}^1 \sqrt{1-x^2}dx &= \int_\pi^0 |\sin\theta|(-\sin\theta)d\theta = \int_0^\pi \sin^2\theta d\theta \\ &= \frac{1}{2}\int_0^\pi (1-\cos(2\theta))d\theta = \frac{1}{2}\left[\theta - \frac{1}{2}\sin(2\theta)\right]_0^\pi \\ &= \frac{\pi}{2}.\end{aligned}$$

So our iterated integral evaluates to $\boxed{\pi/2}$. Notice that this is the area of a semicircular disc with radius 1. If we interpret our bounds of integration as giving the region

$$\{(x,y) : -1 \le x \le 1 \text{ and } 0 \le y \le \sqrt{1-x^2}\}$$

in $\mathbb{R}^2$, we get precisely such a semicircular disc.

□

2.2 Day 5: Fubini's theorem

Goals

By the end of today's class, we should be able to do the following.

1. State **Fubini's theorem**, relating multiple integrals to iterated integrals.
2. Sketch the proof of Fubini's theorem.
3. List the iterated integrals which correspond to a given multiple integral.

2.2.1 Fubini's theorem: from multiple integrals to iterated integrals

We've now gone to great pains to formally define integration for functions of several variables, but we still don't have any truly practical means of evaluating an integral. A similar thing happened in single-variable calculus — at first, the only way you knew how to compute the value of an integral was via Riemann sums, which were tedious and ugly. But then the Fundamental Theorem of Integral Calculus saved the day, telling you that an integral could be computed as an *antiderivative*. This time our savior is Fubini's theorem.

Theorem 2.5: Fubini's Theorem

Let $B = [a,b] \times [c,d]$ be a compact box in $\mathbb{R}^2$, and let $f : B \to \mathbb{R}$ be bounded and nearly-continuous on B. Suppose that for each $x_0 \in [a,b]$, the cross-sectional integral $\int_c^d f(x_0, y)dy$ exists. Then the iterated integral $\int_a^b \int_c^d f(x,y)dydx$ also exists, and

$$\int_B f\,dA = \int_a^b \int_c^d f(x,y)dydx. \tag{2.1}$$

Remark. Some notes about Theorem 2.5:

1. Some version of Fubini's theorem holds in $\mathbb{R}^n$, for any $n \geq 2$. We've just stated the $n = 2$ case here for ease of notation, but you're invited to write the analogous statement in higher dimensions.
2. You should internalize Fubini's theorem as telling us that multiple integrals can be computed as iterated integrals.

Idea of proof. Throughout our proof, $\mathcal{P}_1$ will denote a partition of $[a,b]$, while $\mathcal{P}_2$ is a partition of $[c,d]$, and therefore $\mathcal{P} = \mathcal{P}_1 \times \mathcal{P}_2$ is a partition of B. We'll also consider the function $g : [a,b] \to \mathbb{R}$ defined by

$$g(x) = \int_c^d f(x,y)dy.$$

Based on our discussions of the last several days, we have

$$L(f,\mathcal{P}) \leq \int_B f\,dA \leq U(f,\mathcal{P}) \tag{2.2}$$

for the integral on the left hand side of (2.1), and

$$L(g,\mathcal{P}_1) \leq \int_a^b g(x)dx \leq U(g,\mathcal{P}_1) \tag{2.3}$$

for the integral on the right. Moreover, as we take finer and finer partitions, these two inequalities will squeeze in on the true values of the respective integrals. The key step in our proof will be to argue that we have

$$L(f,\mathcal{P}) \leq L(g,\mathcal{P}_1) \leq U(g,\mathcal{P}_1) \leq U(f,\mathcal{P}). \tag{2.4}$$

Now when the inequalities (2.2) and (2.3) squeeze down on their respective values, we see from (2.4) that these values are in fact the same.

Some details of the argument will be left as (optional) exercises. Throughout, we'll use the notation $\varphi_x(y) := f(x,y)$. Notice that $g(x) = \int_c^d \varphi_x(y)dy$, for every $x \in [a,b]$.

Exercise 2.6. For any subinterval J associated to $\mathcal{P}_1$ and any subinterval K associated to $\mathcal{P}_2$, show that

$$m_{J\times K}(f) \le m_K(\varphi_x),$$

for every $x \in J$.

Using Exercise 2.6, we have the following chain of inequalities:

$$\begin{aligned}\sum_K m_{J\times K}(f)|K| &\le \sum_K m_K(\varphi_x)|K| = L(\varphi_x, \mathcal{P}_2)\\ &\le \int_c^d \varphi_x(y)dy = g(x),\end{aligned}$$

for all $x \in J$. Here J is an arbitrary subinterval associated to $\mathcal{P}_1$, and the sums are taken over subintervals K associated to $\mathcal{P}_2$. The point is that we have a lower bound for g on the subinterval J, and thus

$$\sum_K m_{J\times K}(f)|K| \le m_J(g). \tag{2.5}$$

Exercise 2.7. Use (2.5) to show that

$$\sum_{J,K} m_{J\times K}(f)\,\mathrm{vol}(J\times K) \le \sum_J m_J(g)|J|$$

where the sums are taken over all subintervals J, K associated to $\mathcal{P}_1, \mathcal{P}_2$, respectively.

Finally, we may rewrite the conclusion of Exercise 2.7 as

$$L(f,\mathcal{P}) \le L(g,\mathcal{P}_1),$$

which gives us one half of the key inequality (2.4). A completely analogous inequality tells us that $U(g,\mathcal{P}_1) \le U(f,\mathcal{P})$, and thus we get (2.4). □

2.2.2 The iterated integrals associated to a multiple integral

Now that we've convinced ourselves that Fubini's theorem is true, let's practice using it.

Example 2.8. Evaluate the multiple integral $\int_B f\,dA$, where $B = [0,1]\times[0,2]$ and $f(x,y) = x^2 + y^2$.

Solution. First, let's point out that f is definitely integrable over B, since B is compact and f is continuous[2]. According to Fubini's theorem, we have

$$\begin{aligned}\int_B f\,dA &= \int_0^1 \left(\int_0^2 (x^2+y^2)dy\right)dx = \int_0^1 \left[x^2y + \frac{1}{3}y^3\right]_0^2 dx\\ &= \int_0^1 \left(2x^2 + \frac{8}{3}\right)dx = \left[\frac{2}{3}x^3 + \frac{8}{3}x\right]_0^1\\ &= \boxed{\frac{10}{3}}.\end{aligned}$$

Notice that this agrees with the value of $\int_0^2\int_0^1 (x^2+y^2)dydx$, which we computed in Example 2.4. Indeed, provided the relevant hypotheses are satisfied, Fubini's theorem allows us to compute a multiple integral as an iterated integral *in any order we like*. □

Fubini's theorem is stated for boxes, but it also allows us to compute integrals over nonboxes as iterated integrals. For instance, suppose we have functions $g, h\colon [a,b]\to\mathbb{R}$, with $g(x) \le h(x)$ for all $x \in [a,b]$. Then we can define a region

$$R := \{(x,y) : x \in [a,b] \text{ and } g(x) \le y \le h(x)\} \subset \mathbb{R}^2,$$

which we say is a **vertically simple** region, and then hope to compute $\int_R f\,dA$, for some function $f : R \to \mathbb{R}^2$. Fubini's theorem allows us to compute this integral as an iterated integral.

[2]We won't be in the habit of saying this every time we compute an integral; I'm just pointing it out here since this is our first computation.

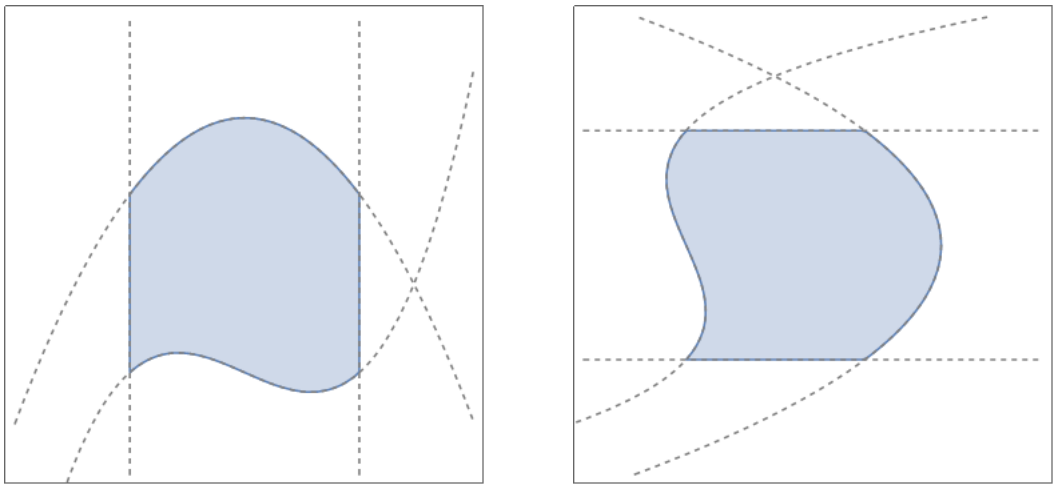

(a) Vertically simple (b) Horizontally simple

Figure 2.1: A vertically simple region and a horizontally simple region in $\mathbb{R}^2$

Exercise 2.9. Use Fubini's theorem to prove that

$$\int_R f\,dA = \int_a^b \left(\int_{g(x)}^{h(x)} f(x,y)dy \right) dx$$

under some hypotheses on f, g, and h. (You should specify the hypotheses.)

Example 2.10. Consider the semicircular disc $R = \{(x,y) : -1 \le x \le 1 \text{ and } 0 \le y \le \sqrt{1-x^2}\}$. The area of this region is given by

$$\text{Area}(R) = \int_R 1dA = \int_{-1}^{1} \int_0^{\sqrt{1-x^2}} 1\,dy\,dx = \frac{\pi}{2},$$

where the final equality follows from our work in Example 2.4.

In a manner completely analogous to Exercise 2.9, one can show that if

$$R = \{(x,y) : c \le y \le d \text{ and } g(y) \le x \le h(y)\}$$

for some functions $g, h\colon [c,d] \to \mathbb{R}$ satisfying $g(y) \le h(y)$, then

$$\int_R f\,dA = \int_c^d \int_{g(y)}^{h(y)} f(x,y)dx\,dy,$$

supposing g, h, and f are suitably nice. Regions which can be written in this manner are **horizontally simple**. Many regions we encounter are both vertically simple *and* horizontally simple, meaning that we can compute integrals over these regions as iterated integrals in two different ways: with x as our inner variable, or with y as our inner variable. Sometimes the order of integration has no practical effect, but in other cases, we can only reasonably compute one of the two iterated integrals.

Example 2.11. Let R be the right triangle in $\mathbb{R}^2$ with vertices at $(0,0)$, $(1,2)$, and $(1,0)$, and define $f : R \to \mathbb{R}$ by $f(x,y) = e^{-x^2}$. Let's compute $\int_R f\,dA$. As depicted in Figure 2.2, R is both vertically simple and horizontally simple: we may write

$$R = \{(x,y) : 0 \le x \le 1 \text{ and } 0 \le y \le 2x\}$$

and

$$R = \{(x,y) : 0 \le y \le 2 \text{ and } y/2 \le x \le 1\}.$$

Fubini's theorem tells us that we can compute $\int_R f\,dA$ as either of the two corresponding iterated integrals, so let's think of R as horizontally simple. In this case, we find that

$$\int_R f\,dA = \int_0^2 \int_{y/2}^1 e^{-x^2} dx\,dy.$$

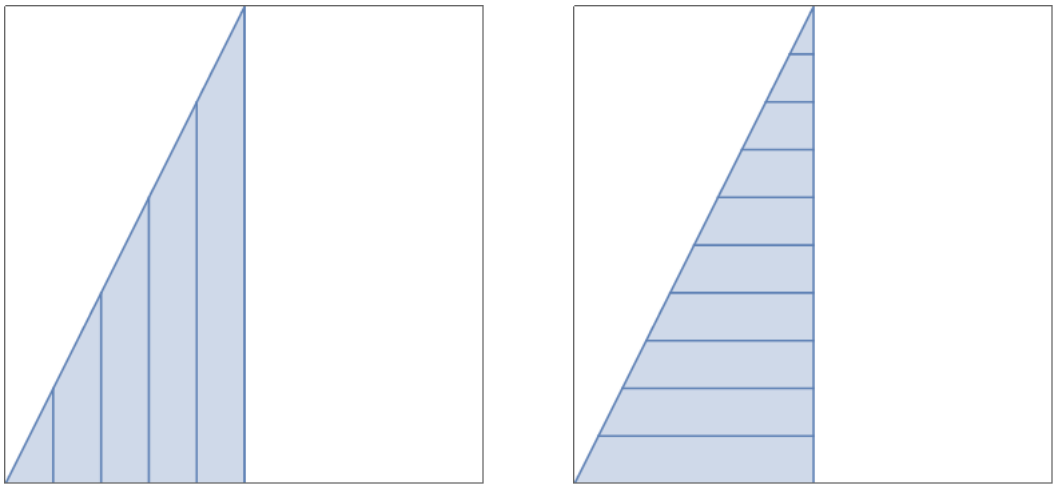

(a) Vertically simple: y varies from 0 to $2x$

(b) Horizontally simple: x varies from $y/2$ to 1

Figure 2.2: The triangle with vertices $(0,0)$, $(1,0)$, and $(1,2)$ is both vertically and horizontally simple.

But this isn't particularly helpful, since e^{-x^2} has no elementary antiderivative. On the other hand, thinking of R as vertically simple leads us to write

$$\int_R f\,dA = \int_0^1 \int_0^{2x} e^{-x^2}\,dy\,dx = \int_0^1 2xe^{-x^2}\,dx = \int_0^1 e^{-u}\,du = 1 - e^{-1}.$$

The penultimate equality is obtained using the substitution $u = x^2$. So while Fubini's theorem tells us that the order of integration is unimportant as an abstract matter, order is often very important as a practical matter.

Because we're all so very talented when it comes to computing single-variable integrals, the most difficult part of computing a multiple integral will often be accurately expressing it as an iterated integral. Let's finish today with an example for a function of three variables.

Example 2.12. Consider the tetrahedron R in $\mathbb{R}^3$ with vertices $(a,0,0)$, $(0,b,0)$, and $(0,0,c)$, as seen in Figure 2.3. Given a sufficiently nice function f, we can use Fubini's theorem to produce six different iterated integrals which compute the value of $\int_R f\,dV$. We get a different iterated integral for each choice of ordering for the variables x, y, and z.

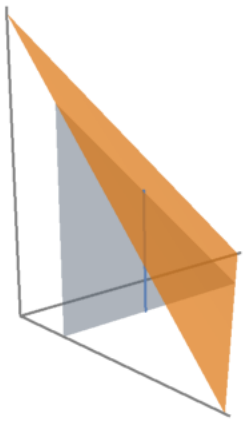

Figure 2.3: A tetrahedron bounded by the (orange) plane $bcx + acy + abz = abc$, as well as the coordinate planes.

For instance, say we want to write

$$\int_R f\,dV = \int_{x_-}^{x^+} \int_{y_-(x)}^{y_+(x)} \int_{z_-(x,y)}^{z_+(x,y)} f(x,y,z)dzdydx,$$

for some functions $y_-(x) \le y_+(x)$ and $z_-(x,y) \le z_+(x,y)$ and some constants $x_- \le x_+$. Notice that the limits of integration for each integral are allowed to depend only on those variables which are integrated "further out." So the innermost integral has bounds depending on both x and y, while the bounds of the outermost integral are constant. We can determine our bounds of integration by working from the outside in. We first ask what range of x-values are seen in our region. This is easy: x is allowed to range from 0 to a, so $x_- = 0$ and $x_+ = a$. Next, for a given x-value, what values may y attain? Consider the $x =$ constant slice seen in Figure 2.3. Certainly y is allowed to be as small as 0 (here we are assuming that $a, b, c > 0$); at the other extreme, y can be as large as $b - \frac{b}{a}x$ (when $z = 0$). Finally, we move to the innermost integral. Here we know the values of x and y and ask the values over which z is allowed to range. This is depicted by the vertical line in Figure 2.3. The smallest value achieved by z is 0, while the largest is on the plane

$$bcx + acy + abz = abc,$$

and thus corresponds to $z = c - \frac{c}{a}x - \frac{c}{b}y$. Altogether, we've found that we can write the tetrahedron R as

$$R = \left\{ (x,y,z) \quad : \quad \begin{array}{c} 0 \le x \le a \\ 0 \le y \le b - (b/a)x \\ 0 \le z \le c - (c/a)x - (c/b)y \end{array} \right\}.$$

From this it follows that

$$\int_R f\,dV = \int_0^a \int_0^{b-(b/a)x} \int_0^{c-(c/a)x-(c/b)y} f(x,y,z)dzdydx.$$

Now say we want (for whatever reason) to integrate against y first, then x, and then z. In order to obtain an iterated integral of this form, we need to write

$$R = \left\{ (x,y,z) \quad : \quad \begin{array}{c} z_- \le z \le z_+ \\ x_-(z) \le x \le x_+(z) \\ y_-(x,z) \le y \le y_+(x,z) \end{array} \right\},$$

for some constants $z_- \le z_+$ and functions $x_-(z) \le x_+(z)$ and $y_-(x,z) \le y_+(x,z)$. As above, we can argue geometrically to write

$$\int_R f\,dV = \int_0^c \int_0^{a-(a/c)z} \int_0^{b-(b/a)x-(b/c)z} f(x,y,z)dydxdz.$$

Exercise 2.13. Express $\int_R f\,dV$ as an iterated integral in four other ways.

2.3 Day 6: Fubini's theorem in action

Goals

By the end of today's class, we should be able to do the following.

1. Use Fubini's theorem to evaluate multiple integrals.
2. Use multiple integrals to compute volumes, centroids, and moments of inertia.

Today will be devoted to examples of using Fubini's theorem to evaluate multiple integrals. In the course of these examples, we will also encounter a few applications of multiple integrals.

Example 2.14. Let W be the region in the first octant of $\mathbb{R}^3$ bounded by the plane $3x + 4y + 2z = 12$, the plane $x + 2y = 3$, and the parabolic cylinder $y^2 = x$. Find the volume of W.

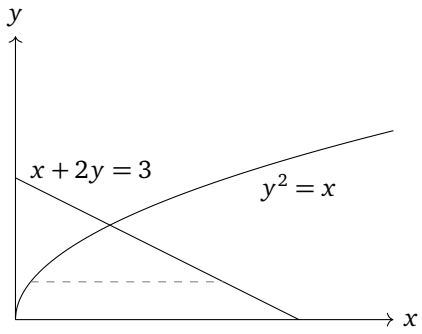

Figure 2.4: The z-independent conditions of Example 2.14.

Solution. Yikes. We have three equalities that we need to turn into inequalities:

$$3x + 4y + 2z = 12, \quad x + 2y = 3, \quad y^2 = x.$$

Additionally, we know that $x, y, z \geq 0$, since W lies in the first octant of $\mathbb{R}^3$. Perhaps the best observation we can start with is that the last two equalities depend only on x and y, so we can think about these in the xy-plane. See Figure 2.4.

Notice that the region D in the first quadrant of $\mathbb{R}^2$ bounded by $y^2 = x$ and $x + 2y = 3$ is horizontally simple. When these two curves meet, we have $y^2 + 2y = 3$, so the maximum value y is allowed to take in D is $y = 2$. Once we've chosen a value for y (as indicated by the dashed gray line in Figure 2.4), x varies from a lower bound of y^2 to an upper bound of $3 - 2y$. Altogether, we may write

$$D = \{(x, y) : 0 \leq y \leq 2, y^2 \leq x \leq 3 - 2y\}.$$

We now have bounds for all but our innermost integral. Namely, we have

$$\int_W f\,dV = \int_0^2 \int_{y^2}^{3-2y} \int_{z_-(x,y)}^{z_+(x,y)} f(x, y, z)\,dz\,dx\,dy,$$

for some $z_-(x, y) \leq z_+(x, y)$. We can determine z_- and z_+ by enforcing the bound $3x + 4y + 2z = 12$. Since z is bounded below by the xy-plane, this plane gives an upper bound. Rearranging, we find that

$$0 \leq z \leq 6 - \frac{3}{2}x - 2y.$$

At last, we can compute the volume of W:

$$\begin{aligned}
\text{vol}(W) &= \int_W 1\,dV = \int_0^2 \int_{y^2}^{3-2y} \int_0^{6-(3/2)x-2y} 1\,dz\,dx\,dy \\
&= \int_0^2 \int_{y^2}^{3-2y} \left[6 - \frac{3}{2}x - 2y\right] dx\,dy \\
&= \int_0^2 \left[6x - \frac{3}{4}x^2 - 2xy\right]_{y^2}^{3-2y} dy \\
&= \int_0^2 \left[\frac{45}{4} - 9y - 5y^2 + 2y^3 + \frac{3}{4}y^4\right] dy \\
&= \left[\frac{45}{4}y - \frac{9}{2}y^2 - \frac{5}{3}y^3 + \frac{1}{2}y^4 + \frac{3}{20}y^5\right]_0^2 \\
&= \boxed{\frac{119}{30}}
\end{aligned}$$

□

Definition. Let $f : R \to \mathbb{R}$ be an integrable function on a compact region R of $\mathbb{R}^n$. Then the **average value of f over R** is defined to be the value

$$\overline{f} = \frac{1}{\text{vol}(R)} \int_R f\, dV.$$

Example 2.15. Suppose that the temperature in degrees Celsius at a point (x, y) on a flat metal plate is $T(x, y) = 5xy + x^2$, where x and y are in meters. Find the average temperature of the box $[0, 2] \times [-2, 2]$.

Solution. The volume (i.e., area) of our box R is $2 \times 4 = 8$, so we need to compute the integral $\int_R T\, dA$. But this is pretty straightforward:

$$\begin{aligned}\int_R T\, dA &= \int_0^2 \int_{-2}^2 (5xy + x^2)\, dy\, dx = \int_0^2 \left[\frac{5}{2}xy^2 + x^2 y\right]_{-2}^2 dx \\ &= \int_0^2 4x^2\, dx = \left[\frac{4}{3}x^3\right]_0^2 \\ &= \frac{32}{3}.\end{aligned}$$

So the average temperature of the box R is $\overline{T} = \frac{1}{8} \cdot \frac{32}{3} = \boxed{\frac{4}{3}}$. □

Definition. Let R be a compact region in $\mathbb{R}^n$. The **centroid** or **barycenter** of R is the arithmetic mean position of all points in R. In coordinates, the centroid is given by $\overline{x} = (\overline{x}_1, \ldots, \overline{x}_n)$, where

$$\boxed{\overline{x}_i = \frac{1}{\text{vol}(R)} \int_R x_i\, dV}$$

for $1 \le i \le n$. If R has density function $\rho : R \to [0, \infty)$, then the **center of mass** of R is the arithmetic mean position of all points, weighted by ρ. The coordinates of the center of mass are given by

$$\boxed{\overline{x}_i = \frac{1}{\text{mass}(R)} \int_R \rho(x_1, \ldots, x_n) x_i\, dV}$$

for $1 \le i \le n$.

Exercise 2.16. Determine conditions on the density function ρ under which the centroid of R is also its center of mass.

Example 2.17. Locate the centroid of the region described by $x^n \le y \le 1$, $0 \le x \le 1$. What is the limiting position of the centroid as $n \to \infty$?

Solution. Let's call the region R_n. Per the description, this region is vertically simple, and its volume is given by

$$\text{vol}(R_n) = \int_0^1 \int_{x^n}^1 1\, dy\, dx = \int_0^1 (1 - x^n)\, dx = \left[x - \frac{1}{n+1}x^{n+1}\right]_0^1 = 1 - \frac{1}{n+1} = \frac{n}{n+1}.$$

The x-coordinate of the centroid is then

$$\begin{aligned}\overline{x} &= \frac{n+1}{n} \int_0^1 \int_{x^n}^1 x\, dy\, dx = \frac{n+1}{n} \int_0^1 (1 - x^n) x\, dx \\ &= \frac{n+1}{n}\left[\frac{1}{2}x^2 - \frac{1}{n+2}x^{n+2}\right]_0^1 = \frac{n+1}{n}\left[\frac{1}{2} - \frac{1}{n+2}\right] \\ &= \frac{n+1}{n} \frac{n}{2(n+2)} = \frac{n+1}{2(n+2)}.\end{aligned}$$

At the same time, the y-coordinate of the centroid is

$$\begin{aligned}\overline{y} &= \frac{n+1}{n}\int_0^1\int_{x^n}^1 y\,dy\,dx = \frac{n+1}{n}\int_0^1\left[\frac{1}{2}y^2\right]_{x^n}^1 dx\\ &= \frac{n+1}{n}\int_0^1\left[\frac{1}{2}-\frac{1}{2}x^{2n}\right]dx = \frac{n+1}{n}\left[\frac{1}{2}x-\frac{1}{2(2n+1)}x^{2n+1}\right]_0^1\\ &= \frac{n+1}{n}\frac{n}{2n+1} = \frac{n+1}{2n+1}.\end{aligned}$$

So the centroid is $\boxed{(\overline{x},\overline{y}) = \left(\frac{n+1}{2(n+2)}, \frac{n+1}{2n+1}\right)}$. As $n\to\infty$, this centroid limits to $\boxed{\left(\frac{1}{2},\frac{1}{2}\right)}$. Notice that as $n\to\infty$, our region tends toward the box $[0,1]\times[0,1]$, so the centroid $(1/2,1/2)$ checks out. □

Example 2.18. Find the mass and center of mass of the cube

$$B := \{(x,y,z) : 0\le x,y,z\le a\},$$

assuming that this cube has density function $\rho(x,y,z) = a-x$, where $a>0$ is some positive number.

Solution. The mass of B is given by

$$\begin{aligned}\text{mass}(B) &= \int_B \rho\,dV = \int_0^a\int_0^a\int_0^a (a-x)\,dz\,dy\,dx = \int_0^a a^2(a-x)\,dx\\ &= a^2\left[ax-\frac{1}{2}x^2\right]_0^a = a^2(a^2-\frac{1}{2}a^2) = \frac{1}{2}a^4.\end{aligned}$$

We can find the center of mass one coordinate at a time:

$$\begin{aligned}\int_B \rho x\,dV &= \int_0^a\int_0^a\int_0^a (a-x)x\,dz\,dy\,dx = \int_0^a a^2(a-x)x\,dx\\ &= a^2\left[\frac{a}{2}x^2-\frac{1}{3}x^3\right]_0^a = a^2\left[\frac{1}{2}a^3-\frac{1}{3}a^3\right] = \frac{1}{6}a^5,\end{aligned}$$

so $\overline{x} = \left(\frac{1}{6}a^5\right)/\left(\frac{1}{2}a^4\right) = \frac{1}{3}a$. Next,

$$\begin{aligned}\int_B \rho y\,dV &= \int_0^a\int_0^a\int_0^a (a-x)y\,dz\,dx\,dy = \int_0^a\int_0^a a(a-x)y\,dx\,dy\\ &= \int_0^a\left[a\left(ax-\frac{1}{2}x^2\right)y\right]_0^a dy = \int_0^a \frac{1}{2}a^3y\,dy = \left[\frac{1}{4}a^3y^2\right]_0^a = \frac{1}{4}a^5,\end{aligned}$$

so $\overline{y} = \left(\frac{1}{4}a^5\right)/\left(\frac{1}{2}a^4\right) = \frac{1}{2}a$. Symmetrically, $\overline{z} = \frac{1}{2}a$. So the center of mass of B is $\boxed{\left(\frac{1}{3}a,\frac{1}{2}a,\frac{1}{2}a\right)}$. □

Example 2.19. Set up (but do not evaluate) an integral which computes the x-coordinate of the centroid of an octant of a solid sphere of radius R centered at the origin in $\mathbb{R}^3$.

Solution. For clarity, let's write the octant as

$$W = \{(x,y,z) : 0\le x,y,z, x^2+y^2+z^2\le R^2\}.$$

Since we need to compute integral over W, we start by rewriting W in a form amenable to integration:

$$W = \left\{(x,y,z) \quad : \quad \begin{array}{c} 0\le x\le R \\ 0\le y\le\sqrt{R^2-x^2} \\ 0\le z\le\sqrt{R^2-x^2-y^2}\end{array}\right\}.$$

Now we could compute $\text{vol}(W)$ as an integral, but that sounds unnecessarily annoying. We know that the volume of a sphere is $\frac{4}{3}\pi R^3$, so $\text{vol}(W) = \frac{1}{6}\pi R^3$. Next, let's set up $\int_W x\,dV$.

$$\int_W x\,dV = \int_0^R\int_0^{\sqrt{R^2-x^2}}\int_0^{\sqrt{R^2-x^2-y^2}} x\,dz\,dy\,dx$$

So

$$\boxed{\overline{x} = \frac{1}{\text{vol}(W)} \int_W x\,dV = \frac{6}{\pi} \int_0^R \int_0^{\sqrt{R^2-x^2}} \int_0^{\sqrt{R^2-x^2-y^2}} x\,dz\,dy\,dx.}$$

This integral would be a pain to evaluate; next week we'll find that it's much easier to evaluate in a different coordinate system. □

Definition. Let R be a compact region in $\mathbb{R}^n$ whose density at $\vec{x} \in R$ is given by $\rho(\vec{x})$, and let L be a straight line in $\mathbb{R}^n$. The **moment of inertia of** R **about** L is given by

$$I_L = \int_R \rho D^2 dV,$$

where $D : R \to [0, \infty)$ measures the distance from L.

Example 2.20. Find the moment of inertia of the solid cylinder

$$C := \{(x, y, z) : x^2 + y^2 \leq a^2, 0 \leq z \leq h\}$$

about the z-axis, where $a, h > 0$ are positive numbers.

Solution. Because we aren't given a density function, we assume that $\rho \equiv 1$. Next, for any point $(x, y, z) \in C$, the distance from (x, y, z) to the z-axis is given by $D(x, y, z) = \sqrt{x^2 + y^2}$. So the relevant moment of inertia is computed as

$$I_z = \int_C (x^2 + y^2) dV.$$

Using Fubini's theorem, we compute this as an iterated integral:

$$\begin{aligned} I_z &= \int_0^h \int_{-a}^a \int_{-\sqrt{a^2-x^2}}^{\sqrt{a^2-x^2}} (x^2 + y^2) dy\,dx\,dz = \int_0^h \int_{-a}^a \left[x^2 y + \frac{1}{3} y^3 \right]_{-\sqrt{a^2-x^2}}^{\sqrt{a^2-x^2}} dx\,dz \\ &= \int_0^h \int_{-a}^a \frac{1}{3} (a^2 - x^2)^{3/2} dx\,dz = \int_{-a}^a \int_0^h \frac{1}{3} (a^2 - x^2)^{3/2} dz\,dx \\ &= \int_{-a}^a \frac{h}{3} (a^2 - x^2)^{3/2} dx. \end{aligned}$$

Notice that we used the penultimate equality to swap the order of integration. You should be able to use your single-variable integration skills to show that this last integral evaluates to $\boxed{\frac{1}{8} a^4 h \pi}$, but it's not very fun. Again, our life will be made easier when we have access to other coordinate systems. □

Example 2.21. For any $n \geq 0$, the **standard** n**-simplex** is the region

$$\Delta^n := \left\{ (t_1, \ldots, t_n) \in \mathbb{R}^n : \sum_{i=1}^n t_i \leq 1 \text{ and } t_1, \ldots, t_n \geq 0 \right\}.$$

Use multiple integration[3] to find a formula for the volume of Δ^n, for all $n \geq 1$.

Solution. We could jump straight to computing $\text{vol}(\Delta^n)$, for an arbitrary n, without too much trouble, but let's start by considering the $n = 1$, $n = 2$, and $n = 3$ cases on their own. First,

$$\Delta^1 = \{ t_1 \in \mathbb{R} : t_1 \leq 1 \text{ and } t_1 \geq 0 \}.$$

Computing the volume here is easy enough: we simply have

$$\text{vol}(\Delta^1) = \int_0^1 1 dt_1 = 1.$$

[3]This is probably not the easiest way to find this volume, but then we didn't sign up for 32BH to do things the easy way.

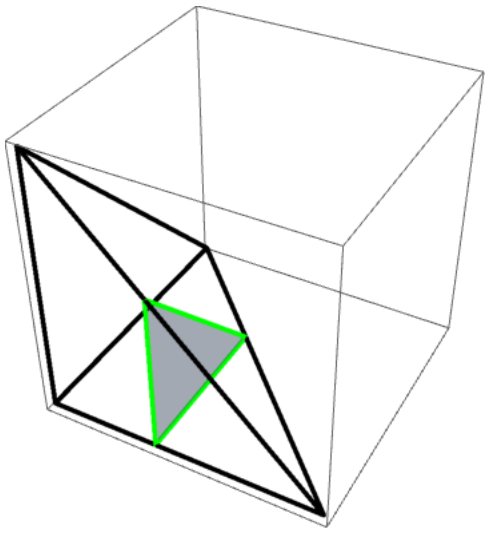

Figure 2.5: Cutting the 3-simplex Δ^3 yields a 2-simplex $t_1\Delta^2$.

Next, there's the 2-simplex

$$\Delta^2 = \{(t_1,t_2)\in\mathbb{R}^2 : t_1+t_2\le 1 \text{ and } t_1,t_2\ge 0\} = \{(t_1,t_2)\in\mathbb{R}^2 : 0\le t_1\le 1, 0\le t_2\le 1-t_1\},$$

which has volume

$$\mathrm{vol}(\Delta^2) = \int_0^1\int_0^{1-t_1} 1\,dt_2\,dt_1 = \int_0^1 (1-t_1)dt_1 = \left[t_1-\frac{1}{2}t_1^2\right]_0^1 = \frac{1}{2}.$$

For our last concrete example, we have the 3-simplex

$$\begin{aligned}\Delta^3 &= \{(t_1,t_2,t_3)\in\mathbb{R}^3 : t_1+t_2+t_3\le 1 \text{ and } t_1,t_2,t_3\ge 0\}\\ &= \{(t_1,t_2,t_3)\in\mathbb{R}^3 : 0\le t_1\le 1, 0\le t_2\le 1-t_1, 0\le t_3\le 1-t_1-t_2\},\end{aligned}$$

so

$$\begin{aligned}\mathrm{vol}(\Delta^3) &= \int_0^1\int_0^{1-t_1}\int_0^{1-t_1-t_2} 1\,dt_3\,dt_2\,dt_1 = \int_0^1\int_0^{1-t_1}(1-t_1-t_2)dt_2\,dt_1\\ &= \int_0^1\left[t_2-t_1t_2-\frac{1}{2}t_2^2\right]_0^{1-t_1}dt_1 = \int_0^1\left[(1-t_1)-t_1(1-t_1)-\frac{1}{2}(1-t_1)^2\right]dt_1\\ &= \int_0^1\frac{1}{2}(t_1-1)^2dt_1 = \left[\frac{1}{6}(t_1-1)^3\right]_0^1 = \frac{1}{6}.\end{aligned}$$

Here's an interesting observation: in the above computation, the outermost integral has t_1 as its variable. If we cut the 3-simplex Δ^3 with the plane $t_1 = c$, where c is some constant, we get a scaled copy $c\Delta^2$ of the standard 2-simplex. See Figure 2.5. Now we know that $\mathrm{vol}(\Delta^2) = \frac{1}{2}$, so we should have $\mathrm{vol}(c\Delta^2) = \frac{1}{2}c^2$. This is because volume in $\mathbb{R}^n$ scales as

$$\mathrm{vol}(cK) = c^n\,\mathrm{vol}(K), \tag{2.6}$$

where K is some compact region.

Exercise 2.22. Prove equation 2.6.

So we can compute $\mathrm{vol}(\Delta^3)$ by integrating the function $\frac{1}{2}t_1^2$ as t_1 varies from 0 to 1:

$$\mathrm{vol}(\Delta^3) = \int_0^1\frac{1}{2}t_1^2dt_1 = \left[\frac{1}{6}t_1^3\right]_0^1 = \frac{1}{6}.$$

Completely analogously, we can find the volume of Δ^4:

$$\text{vol}(\Delta^4) = \int_0^1 \frac{1}{6} t_1^3 dt_1 = \left[\frac{1}{24} t_1^4\right]_0^1 = \frac{1}{24}.$$

Perhaps by now you've noticed the pattern $\text{vol}(\Delta^n) = \frac{1}{n!}$. We can prove this by mathematical induction. We already have several base cases, so let's assume that $\text{vol}(\Delta^n) = \frac{1}{n!}$ for some $n \geq 1$ and show that $\text{vol}(\Delta^{n+1}) = \frac{1}{(n+1)!}$. Indeed,

$$\text{vol}(\Delta^{n+1}) = \int_0^1 \frac{1}{n!} t_1^n dt_1 = \left[\frac{1}{(n+1)!} t_1^{n+1}\right]_0^1 = \frac{1}{(n+1)!},$$

verifying our formula $\boxed{\text{vol}(\Delta^n) = \frac{1}{n!}}$. □

3 Week 3

Now that we have Fubini's theorem, we can rewrite our multiple integrals as iterated integrals, and then evaluate. But when the region over which we're integrating is not a box, our bounds of integration can get messy. We saw this in some examples last week, when our domains of integration were cylindrical or spherical. Just as discs in $\mathbb{R}^2$ are more easily written in polar coordinates, we want this week to introduce coordinate systems on $\mathbb{R}^3$ — cylindrical coordinates and spherical coordinates — in which we may more easily express cylinders and spheres. Our real goal is to integrate in these coordinate systems, so the main feature of this week is the change of variables theorem, which will allow us to compute integrals in any coordinate system we can cook up.

3.1 Day 7: Change of variables

Goals

By the end of today's class, we should be able to do the following.

1. Explain how multiple integrals can be **pulled back** via a C^1-transformation $\Phi\colon \mathbb{R}^n \to \mathbb{R}^n$.
2. Give a geometric sketch of why the **change of variables theorem** should be true for double integrals.
3. Use the change of variables theorem and **polar coordinates** to evaluate double integrals.

3.1.1 Integration in polar coordinates

When defining the Riemann integral of a function $f : D \to \mathbb{R}$ of several variables, we partitioned the domain $D \subset \mathbb{R}^n$ into *subboxes* — subsets of the form

$$\{(x_1, \dots, x_n) : a_i \le x_i \le b_i, \text{ for } 1 \le i \le n\}.$$

For instance, in $\mathbb{R}^2$ a box is simply a rectangle. But this is because we're using *rectangular coordinates*. There are other coordinate systems we could use on $\mathbb{R}^2$, and boxes in these coordinate systems will not necessarily be rectangles. For instance, recall *polar coordinates* on $\mathbb{R}^2$. Each point (x, y) in $\mathbb{R}^2$ can be written as

$$(x, y) = (r\cos\theta, r\sin\theta)$$

for some $r \in \mathbb{R}_{\ge 0}$ and $\theta \in \mathbb{R}$. (Of course, r and θ are not uniquely defined, but let's not worry about this just yet.) A *polar box* then has the form

$$\{(r\cos\theta, r\sin\theta) : r_1 \le r \le r_2, \theta_1 \le \theta \le \theta_2\},$$

for some constants[1] $r_2 > r_1 \ge 0$ and $\theta_2 > \theta_1$. If we wanted to define the Riemann integral using polar coordinates, we would think about partitioning a domain D into polar boxes, and then consider lower and upper sums as before:

$$L(f, \mathcal{P}) := \sum_J m_J(f)\text{area}(J) \quad \text{and} \quad U(f, \mathcal{P}) := \sum_J M_J(f)\text{area}(J).$$

The only difference here is that, since each J is a polar box, its area is somewhat more annoying to compute than in the rectangular case.

Before we proceed, let's make it clear why thinking about the Riemann integral in polar coordinates is something worth doing. To this end, consider the unit disc D in $\mathbb{R}^2$:

$$D = \{(x, y) : x^2 + y^2 \le 1\}.$$

This region is both vertically simple and horizontally simple, so we can write it as

$$D = \{(x, y) : -1 \le x \le 1, -\sqrt{1-x^2} \le y \le \sqrt{1-x^2}\}$$

[1] I don't really care to think about strange cases here, so we assume that $r_2 > r_1$ and $\theta_2 > \theta_1$. We could allow $r_2 \le r_1$ and $\theta_2 \le \theta_1$, so that a polar box could have area zero (or even be empty), but let's not worry about this.

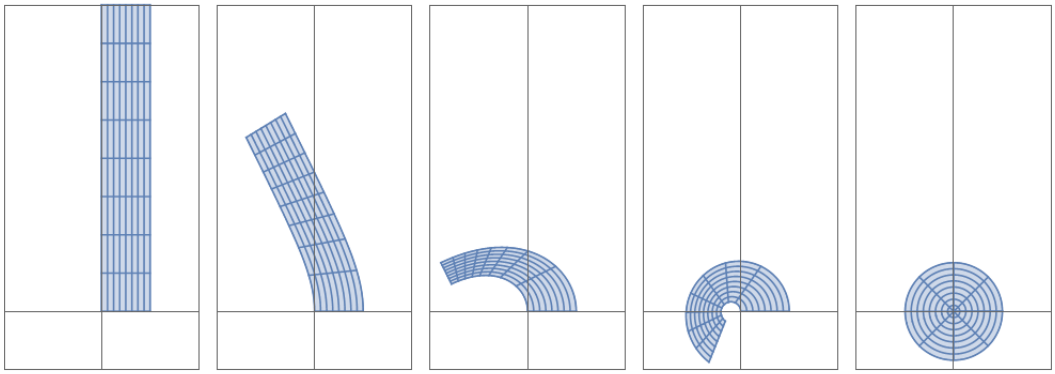

Figure 3.1: As the rectangle $[0,1]\times[0,2\pi]$ is transformed into the unit disc, its subboxes are distorted.

or

$$D=\{(x,y):-1\le y\le 1,-\sqrt{1-y^2}\le x\le\sqrt{1-y^2}\}.$$

We can then use Fubini's theorem to compute a double integral over D as an iterated integral:

$$\int_D f\ dA=\int_{-1}^{1}\int_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} f(x,y)\,dy\,dx=\int_{-1}^{1}\int_{-\sqrt{1-y^2}}^{\sqrt{1-y^2}} f(x,y)\,dx\,dy.$$

Of course, the term $\sqrt{1-x^2}$ or $\sqrt{1-y^2}$ could make our iterated integral computation less than pleasant. But notice that D is a polar rectangle:

$$D=\{(r\cos\theta,r\sin\theta):0\le r\le 1,0\le\theta\le 2\pi\}.$$

So, assuming some sort of Fubini's theorem holds for Riemann integration in polar coordinates, it seems like $\int_D f\ dA$ *ought* to be

$$\int_0^{2\pi}\int_0^1 f(r\cos\theta,r\sin\theta)\,dA.$$

Notice that we've written dA instead of $drd\theta$; this is because we haven't yet sorted out the issue of the area of a polar box. A more onerous example is given by a quarter of an annulus:

$$R=\{(x,y):1\le x^2+y^2\le 4,y\ge|x|\}.$$

This region is neither vertically nor horizontally simple (you should plot the region), so would be pretty annoying to integrate over in rectangular coordinates. But in polar coordinates we have

$$R=\{(r\cos\theta,r\sin\theta):1\le r\le 2,\pi/4\le\theta\le 3\pi/4\},$$

meaning that R is a polar rectangle. Our main motivation for integrating in non-rectangular coordinate systems is that doing so will allow us to treat non-rectangular regions as rectangles — making for iterated integrals which are easier to compute.

3.1.2 Coordinate systems

Based on the informal discussion above, it sounds like we have a shot at interpreting the Riemann integral in polar coordinates if we just compute the area of a polar rectangle. But, being mathematicians, we'd really like to solve this problem for *any* coordinate system. Let's first say what we mean by *coordinate system*, a notion we will formalize via maps $\Phi\colon \mathbb{R}^n \to \mathbb{R}^n$. The definition of coordinate system requires some background vocabulary first.

Definition. We say that a differentiable map $\Phi\colon A \to \mathbb{R}^n$, defined on a subset A of $\mathbb{R}^n$ is a **$\mathbf{C^1}$-mapping** if all of its partial derivatives are continuous. We denote by Φ_* the $n \times n$ matrix of all first-order partial derivatives of Φ:

$$\Phi_* := \begin{pmatrix} \frac{\partial \phi_1}{\partial x_1} & \cdots & \frac{\partial \phi_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \frac{\partial \phi_n}{\partial x_1} & \cdots & \frac{\partial \phi_n}{\partial x_n} \end{pmatrix},$$

where

$$\Phi(x_1, \ldots, x_n) = (\phi_1(x_1, \ldots, x_n), \ldots, \phi_n(x_1, \ldots, x_n)).$$

We call Φ_* the **Jacobian matrix**[2] associated to Φ.

Example 3.1. We can define a differentiable map $\Phi\colon \mathbb{R}^2 \to \mathbb{R}^2$ via

$$\Phi(r, \theta) := (r\cos\theta, r\sin\theta).$$

Then

$$\Phi_* = \begin{pmatrix} \cos\theta & -r\sin\theta \\ \sin\theta & r\cos\theta \end{pmatrix}.$$

Notice that Φ_* is a function of r and θ. To be clear about which variables represent which copy of $\mathbb{R}^2$, we will often denote this map by $\Phi\colon \mathbb{R}^2_{r,\theta} \to \mathbb{R}^2_{x,y}$.

We need a little more vocabulary before defining coordinate systems.

Definition. The **determinant** of a 2×2 matrix is given by the formula

$$\det \begin{pmatrix} a & b \\ c & d \end{pmatrix} = ad - bc.$$

For $n \geq 3$, we may define the determinant of an $n \times n$ matrix recursively:

$$\det \begin{pmatrix} a_{11} & \cdots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{n1} & \cdots & a_{nn} \end{pmatrix} = \sum_{i=1}^{n} (-1)^{i+n} a_{in} \det \begin{pmatrix} a_{11} & \cdots & a_{1(n-1)} \\ \vdots & \cdots & \vdots \\ \widehat{a_{i1}} & \ddots & \widehat{a_{i(n-1)}} \\ \vdots & \cdots & \vdots \\ a_{n1} & \cdots & a_{n(n-1)} \end{pmatrix},$$

where the notation $\widehat{a_{ij}}$ means we omit the entry a_{ij}.

Remark. This is not my favorite way to define the determinant, but it will suffice for now. Today we'll focus fairly exclusively on the case $n = 2$, and hopefully you recall the $n = 3$ case from computing cross products.

Exercise 3.2. Prove that

$$\det \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

gives the signed area of the parallelogram in $\mathbb{R}^2$ spanned by the vectors $(a, c)^T$ and $(b, d)^T$.

[2] Sometimes I might slip and just call Φ_* the *Jacobian* of Φ. I'll try to call this the Jacobian matrix, though, because a lot of textbooks reserve the name 'Jacobian' for the determinant of Φ_*.

Definition. If $\Phi\colon A \to \mathbb{R}^n$ is a C^1-mapping defined on a subset A of $\mathbb{R}^n$, then we call the determinant $\det \Phi_*$ the **Jacobian determinant**[3] of Φ.

Alright, one more simple definition before we say what we mean by coordinate system.

Definition. Let $B \subset \mathbb{R}^n$ be a (not-necessarily-compact) box of the form

$$B = I_1 \times I_2 \times \cdots \times I_n,$$

where each I_i is an interval of one of the following forms:

$$I_i = (a_i, b_i) \quad \text{or} \quad I_i = (a_i, b_i] \quad \text{or} \quad I_i = [a_i, b_i) \quad \text{or} \quad I_i = [a_i, b_i].$$

The **interior** of I_i is $\mathring{I}_i := (a_i, b_i)$, and the interior of B is

$$\mathring{B} := \mathring{I}_1 \times \mathring{I}_2 \times \cdots \times \mathring{I}_n.$$

Note: We don't require the intervals I_i to be bounded, so, read correctly, this definition covers cases like $I_i = (-\infty, b_i]$. Just don't write something silly like $[-\infty, b_i)$.

A coordinate system will me a map whose domain is a box, and we require the map to behave particularly nicely on the interior of the box. The basic idea is that the boundary (i.e., the non-interior points) doesn't have any volume, so if Φ behaves badly there, it shouldn't affect the definition of the integral.

Definition. Given an open subset $A \subset \mathbb{R}^n$, a **coordinate system** on A is a C^1-mapping $\Phi\colon R \to \mathbb{R}^n$ such that

(1) R is a box in $\mathbb{R}^n$;

(2) $A \subset \Phi(R)$;

(3) Φ is injective on the interior of R;

(4) $\det \Phi_* \neq 0$ on the interior of R.

Example 3.3. The map $\Phi\colon \mathbb{R}^2_{r,\theta} \to \mathbb{R}^2_{x,y}$ defined above is not a coordinate system, because it's not injective — in fact, every point in $\mathbb{R}^2_{x,y}$ has infinitely many preimages under Φ. However, the restriction

$$\Phi\colon [0,\infty)_r \times [0, 2\pi]_\theta \to \mathbb{R}^2_{x,y}$$

is a coordinate system on $\mathbb{R}^2$. Indeed:

(1) $[0,\infty) \times [0, 2\pi]$ is a box;

(2) $\mathbb{R}^2 \subset \Phi([0,\infty) \times [0, 2\pi])$;

(3) the restriction $\Phi\colon (0,\infty) \times (0, 2\pi) \to \mathbb{R}^2$ is injective;

(4) we have

$$\det \Phi_* = \det \begin{pmatrix} \cos\theta & -r\sin\theta \\ \sin\theta & r\cos\theta \end{pmatrix} = r\cos^2\theta + r\sin^2\theta = r,$$

which is nonzero on the interior of $(0,\infty) \times (0, 2\pi)$.

Example 3.4. The map $\Phi\colon [1, 3]_u \times [-1, 1]_v \to \mathbb{R}^2_{x,y}$ defined by

$$\Phi(u, v) := (u\cosh(v/2), u\sinh(v/2))$$

is a coordinate system on its image $\Phi(R)$. Recall that the hyperbolic sine and cosine functions satisfy $\cosh^2 v - \sinh^2 v = 1$, and

$$\frac{d}{dv}\cosh v = \sinh v, \qquad \frac{d}{dv}\sinh v = \cosh v.$$

One can check that Φ is injective on all of $R = [1, 3] \times [-1, 1]$ (not just its interior), and we also have

$$\det \Phi_* = \det \begin{pmatrix} \cosh(v/2) & \frac{1}{2}u\sinh(v/2) \\ \sinh(v/2) & \frac{1}{2}u\cosh(v/2) \end{pmatrix} = \frac{1}{2}u\cosh^2(v/2) - \frac{1}{2}u\sinh^2(v/2) = \frac{1}{2}u.$$

[3] Many textbooks just call $\det \Phi_*$ the *Jacobian* of Φ, but we'll call it the Jacobian determinant, to distinguish it from the matrix Φ_*.

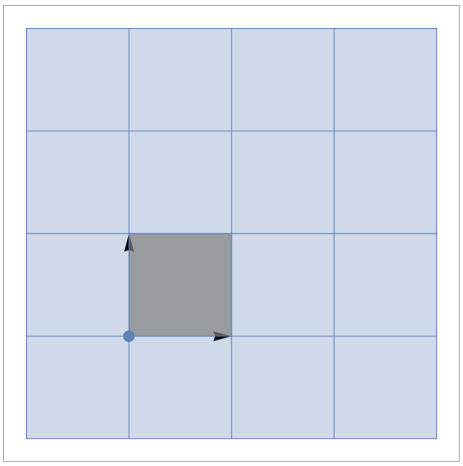

(a) A $\Delta u \times \Delta v$ rectangle based at (u_0, v_0).

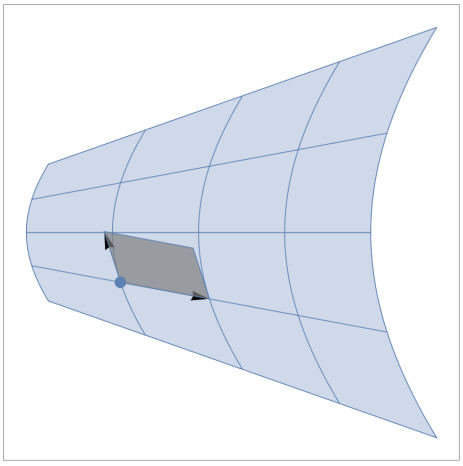

(b) The parallelogram spanned by $\Delta u \frac{\partial \Phi}{\partial u}$ and $\Delta v \frac{\partial \Phi}{\partial v}$.

Figure 3.2: The transformation $\Phi : [a,b]_u \times [c,d]_v \to \mathbb{R}^2_{x,y}$ gives a coordinate system on its image.

So $\det \Phi_* \neq 0$ on R, and Φ is a coordinate system on its image. This coordinate system is depicted in Figure 3.2.

3.1.3 Change of variables

Now that we have a more general notion of coordinate system, we can attempt to define integration for any coordinate system on (a subset of) $\mathbb{R}^n$. As discussed in Section 3.1.1, we may partition the region R over which we want to integrate into **pseudoboxes**[4]

$$J := \{\Phi(u^1, \dots, u^n) : u_0^i < u^i < u_1^i, 1 \leq i \leq n\},$$

and then define the lower and upper sums over our partition to be

$$L(f, \mathcal{P}) := \sum_J m_J(f)\text{vol}(J) \quad \text{and} \quad U(f, \mathcal{P}) := \sum_J M_J(f)\text{vol}(J).$$

The primary difficulty, then, is determining the volumes $\text{vol}(J)$. A key difference between a pseudobox and a standard box in $\mathbb{R}^n$ is that the volume of a pseudobox is not a function of its side lengths alone. For instance, the polar rectangles

$$R_1 := \{(r\cos\theta, r\sin\theta) : 0 < r < 1/4, 0 < \theta < \pi/4\}$$

and

$$R_2 := \{(r\cos\theta, r\sin\theta) : 3/4 < r < 1, 0 < \theta < \pi/4\}$$

each have r-side length $1/4$ and θ-side length $\pi/4$, but $\text{area}(R_1) = \frac{\pi}{128}$, while $\text{area}(R_2) = \frac{7\pi}{128}$ — a seven-fold difference in area!

Rather than attempt to compute the volume of a pseudobox exactly — something that's easy enough to do in polar coordinates, but which we don't want to do for an arbitrary Φ — we will obtain a first-order approximation for this volume. Heuristically, this approximation should be good enough to compute integrals in our new coordinate system, because as the side lengths of our pseudobox shrink to 0, the error in our approximation will tend to 0 even more quickly. That is, if our side lengths are on the order of $\epsilon > 0$, then the error in our volume approximation should be on the order of ϵ^2.

Let's consider the pseudobox J which is the image of $[u_0, u_0 + \Delta u] \times [v_0, v_0 + \Delta v]$. Our first approximation to J is by a quadrilateral: since J has corners at

$$\Phi(u_0, v_0), \quad \Phi(u_0 + \Delta u, v_0), \quad \Phi(u_0, v_0 + \Delta v), \quad \text{and} \quad \Phi(u_0 + \Delta u, v_0 + \Delta v),$$

[4] As far as I'm aware, this language is not at all common. In fact, I just made it up.

we approximate J by the quadrilateral Q with these four corners. We don't give a formal justification for this approximation, but notice that this approximation is accurate for pseudoboxes with straight line edges — making it a reasonable first-order approximation for a pseudobox whose edges are smooth curves. Now computing the area of Q is something we're still too lazy to do, so let's approximate Q by a parallelogram. Namely, we can use Taylor polynomials to approximate three of the four vertices of Q:

$$\begin{aligned}
\Phi(u_0+\Delta u, v_0) &= \Phi(u_0,v_0) + \Delta u \frac{\partial \Phi}{\partial u}(u_0,v_0) + \frac{1}{2}(\Delta u)^2 \frac{\partial^2 \Phi}{\partial u^2}(u_0,v_0) + \cdots \\
&\approx \Phi(u_0,v_0) + \Delta u \frac{\partial \Phi}{\partial u}(u_0,v_0) \\
\Phi(u_0, v_0+\Delta v) &= \Phi(u_0,v_0) + \Delta v \frac{\partial \Phi}{\partial v}(u_0,v_0) + \frac{1}{2}(\Delta v)^2 \frac{\partial^2 \Phi}{\partial v^2}(u_0,v_0) + \cdots \\
&\approx \Phi(u_0,v_0) + \Delta v \frac{\partial \Phi}{\partial v}(u_0,v_0) \\
\Phi(u_0+\Delta u, v_0+\Delta v) &= \Phi(u_0,v_0) + \Delta u \frac{\partial \Phi}{\partial u}(u_0,v_0) + \Delta v \frac{\partial \Phi}{\partial v}(u_0,v_0) + \cdots \\
&\approx \Phi(u_0,v_0) + \Delta u \frac{\partial \Phi}{\partial u}(u_0,v_0) + \Delta v \frac{\partial \Phi}{\partial v}(u_0,v_0).
\end{aligned}$$

With these approximations, Q is approximated by the parallelogram spanned by the vectors $\frac{\partial \Phi}{\partial u}(u_0,v_0)$ and $\frac{\partial \Phi}{\partial v}(u_0,v_0)$ based at (u_0,v_0). See Figure 3.2.

We can now use Exercise 3.2 to estimate the area of a pseudobox with side lengths Δu and Δv. The area of the pseudobox J which is the image under Φ of $[u_0, u_0+\Delta u] \times [v_0, v_0+\Delta v]$ is approximately

$$\boxed{\text{area}(J) \approx |\det \Phi_*(u_0,v_0)|\, \Delta u \Delta v},$$

provided $\Delta u > 0$ and $\Delta v > 0$ are small. By similar reasoning we may obtain the analogous approximation in any dimension. That is, if $\Phi \colon \mathbb{R}^2_{u^1,\ldots,u^n} \to \mathbb{R}^2_{x^1,\ldots,x^n}$ is a coordinate system, then

$$\boxed{\text{vol}(\Phi([u_0^1, u_0^1+\Delta u^1] \times \cdots \times [u_0^n, u_0^n + \Delta u^n])) \approx \left|\det \Phi_*(u_0^1,\ldots,u_0^n)\right| |\Delta u^1 \cdots \Delta u^n|},$$

provided $\Delta u^i > 0$ is small, for $1 \le i \le n$. So the quantity $|\det \Phi_*|$ gives some sort of "fudge factor" for our volume computations.

Where does this leave us? Well, assuming that $|\det \Phi_*|$ doesn't vary too much over each pseudobox — and with a very fine partition it shouldn't — we can approximate our lower sum as

$$\begin{aligned}
L(f,\mathcal{P}) &= \sum_J m_J(f) \text{vol}(J) \approx \sum_J m_J(f) |\det \Phi_*| \Delta u^1 \cdots \Delta u^n \\
&\approx \sum_J m_J(f \cdot |\det \Phi_*|) \text{vol}([u_0^1, u_0^1+\Delta u^1] \times \cdots \times [u_0^n, u_0^n+\Delta u^n])
\end{aligned}$$

and obtain a similar approximation for the upper sum. Taking the supremum over all lower sums and the infimum over all upper sums, we *should* have

$$\int_R f \, dV = \int_D (f \circ \Phi) |\det \Phi_*| \, dV.$$

The following result, which we will not rigorously prove, verifies that this is the case.

Theorem 3.5: The change of variables theorem

Let $K \subset \mathbb{R}^n$ be a compact and connected set having boundary of volume zero. Let $A \subset \mathbb{R}^n$ be an open superset of K, and let $\Phi \colon A \to \mathbb{R}^n$ be a coordinate system on its image. Let $f : \Phi(K) \to \mathbb{R}$ be a continuous function. Then

$$\int_{\Phi(K)} f \, dV = \int_K (f \circ \Phi) |\det \Phi_*| \, dV.$$

Remark.

1. Because $\Phi(K)$ is compact and $f : \Phi(K) \to \mathbb{R}$ is continuous, Theorem 1.23 guarantees that f is integrable over $\Phi(K)$ under the hypotheses of Theorem 3.5.
2. On the left, dV indicates that we're integrating against volume in the codomain of Φ, and on the right dV indicates that we're integrating against volume in the domain of Φ. Theorem 3.5 is telling us that these two "volume forms" are related to each other by the scale factor $\det \Phi_*$. We will attempt to make this more precise later in the quarter.
3. Notice that the domain of f lies in the image of Φ. A typical use of Theorem 3.5 begins with a function $f : \mathbb{R}^2_{x,y} \to \mathbb{R}$ and produces a coordinate system $\Phi : \mathbb{R}^2_{u,v} \to \mathbb{R}^2_{x,y}$. We then use Φ to "pull the integral back" from a complicated region in the xy-plane to a rectangle or other simple region in the uv-plane.

3.1.4 Back to polar coordinates

Remember how we started today: searching for a way to integrate with respect to polar coordinates. Now that we know how to integrate in an arbitrary coordinate system, we can recover the case of polar coordinates. Theorem 3.5 tells us that if Φ denotes the polar coordinate system on $\mathbb{R}^2$ and $R \subset [0, \infty) \times [0, 2\pi]$ is a region in the $r\theta$-plane, then

$$\int_{\Phi(R)} f \, dA = \int_R (f \circ \Phi) r \, dr \, d\theta, \tag{3.1}$$

since Example 3.3 tells us that $|\det \Phi_*| = r$. Let's use Equation 3.1 to compute a value we already know: the area of a semicircular disc.

Remark. I now realize that I've been pretty inconsistent about whether we use R to denote a region in the domain or codomain of a coordinate system. I think I'm sometimes thinking of R as standing for "range," and thus putting it in the codomain of Φ. But other times I'm thinking of R as standing for "rectangle," and putting it in the domain. Hopefully this is not too confusing.

Example 3.6. Let's compute the area of the semicircular disc D of radius 1 centered at the origin. That is,

$$D = \{(x, y) : x^2 + y^2 \leq 1, y \geq 0\} = \{(r\cos\theta, r\sin\theta) : 0 \leq r \leq 1, 0 \leq \theta \leq \pi\},$$

so D is a polar rectangle. Concretely, D is the image under Φ of the rectangle

$$R = \{(r, \theta) : 0 \leq r \leq 1, 0 \leq \theta \leq \pi\} \subset \mathbb{R}^2_{r,\theta}.$$

From this it follows that

$$\begin{aligned} \text{area}(D) &= \int_D 1 \, dA = \int_R r \, dA = \int_0^\pi \int_0^1 r \, dr \, d\theta \\ &= \int_0^\pi \left[\frac{1}{2} r^2\right]_0^1 d\theta = \int_0^\pi \frac{1}{2} \, d\theta = \boxed{\frac{\pi}{2}}, \end{aligned}$$

as expected.

Remark. In a lot of multivariable calculus classes it's popular to think of Equation 3.1 as telling us that "$dA = r dr d\theta$." This is safe enough, provided we interpret the two sides of the equation correctly. Here dA refers to area in $\mathbb{R}^2_{x,y}$. In the $r\theta$-plane, we have $dA_{r,\theta} = dr d\theta$, and we pick up the scale factor of r when we move over to the xy-plane. So maybe a better way to think of this is $dA_{x,y} = r dr d\theta$. Again, we'll attempt to give more precise meaning to all these symbols later in the term.

Here's one more example, which we probably won't have time to see in class.

Example 3.7. Let's find the centroid of the unit-radius semicircular disc D considered above. First, let's note that, even though we're going to use a coordinate change to do this, the coordinates of the centroid are given by computing the average values of the rectangular coordinates of points in D. That is,

$$\overline{x} := \int_D x \, dA \quad \text{and} \quad \overline{y} := \int_D y \, dA.$$

Our coordinate system Φ need not preserve centroids. So we can't expect $\Phi(\overline{r},\overline{\theta})$ to be $(\overline{x},\overline{y})$. Understanding this, let's compute. We have

$$\begin{aligned}\overline{x} &= \int_D x\ dA = \int_R (r\cos\theta) r\ dA = \int_0^\pi \int_0^1 r^2\cos\theta\ dr\ d\theta \\ &= \int_0^\pi \left[\frac{1}{3}r^3\cos\theta\right]_0^1 d\theta = \int_0^\pi \frac{1}{3}\cos\theta\ d\theta = \frac{1}{3}[\sin\theta]_0^\pi = 0.\end{aligned}$$

Next,

$$\begin{aligned}\overline{y} &= \int_D y\ dA = \int_R (r\sin\theta) r\ dA = \int_0^\pi \int_0^1 r^2\sin\theta\ dr\ d\theta \\ &= \int_0^\pi \left[\frac{1}{3}r^3\sin\theta\right]_0^1 d\theta = \int_0^\pi \frac{1}{3}\sin\theta\ d\theta = \frac{1}{3}[-\cos\theta]_0^\pi = \frac{2}{3}.\end{aligned}$$

So the centroid of D has coordinates $(0, \frac{2}{3})$.

3.2 Day 8: Change of variables, continued

Goals

By the end of today's class, we should be able to do the following.

1. Use **cylindrical coordinates** and **spherical coordinates** to evaluate triple integrals.
2. Use the change of variables theorem with wisely-chosen transformations to evaluate multiple integrals in any dimension.
3. In some cases, use the change of variables theorem with incomplete information. For instance, apply the change of variables theorem when we have an expression for the inverse of a coordinate system, but not for the coordinate system itself.

The plan for today is to see several examples of the change of variables theorem in action. In particular, we will see two particularly important coordinate systems for $\mathbb{R}^3$: **cylindrical coordinates** and **spherical coordinates**. Time permitting, we will discuss a situation in which we can apply the change of variables theorem without an explicit expression for our coordinate system (unfortunately, time probably won't permit).

3.2.1 Cylindrical coordinates

We start with an easy generalization of polar coordinates to $\mathbb{R}^3$: cylindrical coordinates. We write a point (x,y,z) in this coordinate system by writing the point (x,y) in polar coordinates. That is,

$$\Phi(r,\theta,z) = (r\cos\theta, r\sin\theta, z).$$

As is the case for polar coordinates, r measures the distance between (x,y,z) and the z-axis (i.e., the distance between (x,y) and the origin in the xy-plane), while θ measures the angle made by the vector $(x,y)^T$ and the x-axis. This coordinate system is depicted in Figure 3.3a.

In order to integrate in cylindrical coordinates, we need to compute the Jacobian determinant. First, the Jacobian matrix is

$$\Phi_* = \begin{pmatrix}\cos\theta & -r\sin\theta & 0\\ \sin\theta & r\cos\theta & 0\\ 0 & 0 & 1\end{pmatrix}.$$

We compute the Jacobian determinant using cofactor expansion:

$$\det\Phi_* = \det\begin{pmatrix}\cos\theta & -r\sin\theta & 0\\ \sin\theta & r\cos\theta & 0\\ 0 & 0 & 1\end{pmatrix} = \det\begin{pmatrix}\cos\theta & -r\sin\theta\\ \sin\theta & r\cos\theta\end{pmatrix} = r.$$

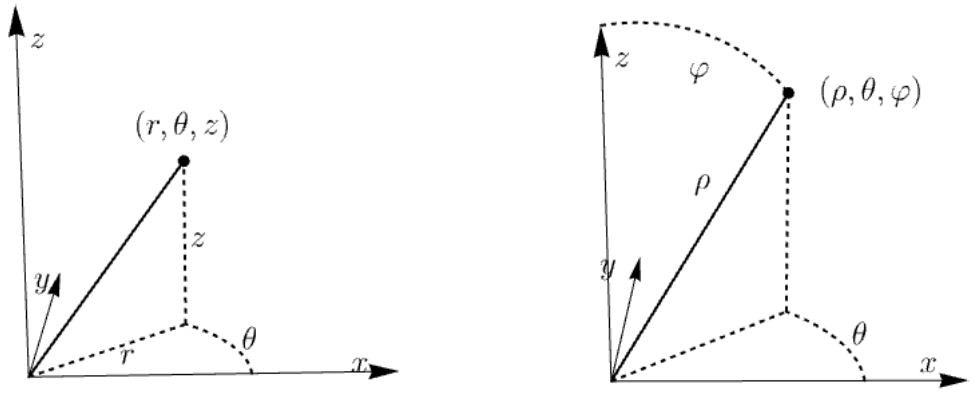

(a) Cylindrical coordinates

(b) Spherical coordinates

Figure 3.3: Polar coordinates generalize to cylindrical and spherical coordinates on $\mathbb{R}^3$

Of course this matches the Jacobian determinant for polar coordinates — not too surprising, since the cylindrical coordinate system is just the result of crossing polar coordinates with z. So the change of variables theorem tells us that

$$\boxed{\int_{\Phi(R)} f\ dV = \int_R (f\circ\Phi)|\det \Phi_*|\ dV = \int_R (f\circ\Phi) r\ dV}.$$

Example 3.8. Let W be the region in $\mathbb{R}^3$ bounded between the paraboloids $z = x^2+y^2-4$ and $z = 4-x^2-y^2$. Find the moment of inertia of W about the z-axis, assuming that W has uniform density ρ.

Solution. Here we have

$$W = \{(x,y,z) : x^2+y^2-4 \le z \le 4-x^2-y^2\},$$

and it should be clear that this region looks quite a bit nicer in cylindrical coordinates:

$$W = \{(r\cos\theta, r\sin\theta, z) : r^2-4 \le z \le 4-r^2\}.$$

Notice that W places no restriction on θ, so we'll let θ vary from 0 to 2π. For the restriction on z to make sense, we'll need $r^2-4 \le 4-r^2$, which occurs only if $r^2 \le 2$. So we'll let r vary from 0 to 2. Now it shouldn't be too bad to integrate over W. In particular,

$$I_z = \int_W \rho\cdot D^2\ dV = \int_0^{2\pi}\int_0^2\int_{r^2-4}^{4-r^2} \rho\cdot D^2\circ\Phi r\ dz\ dr\ d\theta.$$

Finally, we notice that r measures the distance between a point and the z-axis, so $(D\circ\Phi)(r,\theta,z) = r$, and we have

$$\begin{aligned} I_z &= \int_0^{2\pi}\int_0^2\int_{r^2-4}^{4-r^2} \rho\cdot r^3\ dz\ dr\ d\theta = \int_0^{2\pi}\int_0^2 \left[\rho r^3 z\right]_{r^2-4}^{4-r^2}\ dr\ d\theta \\ &= \int_0^{2\pi}\int_0^2 \rho r^3(8-2r^2)\ dr\ d\theta = \int_0^{2\pi}\left[\rho\left(2r^4-\frac{2}{6}r^6\right)\right]_0^2\ d\theta \\ &= \int_0^{2\pi}\rho\left(32-\frac{64}{3}\right)\ d\theta = \boxed{\frac{64}{3}\pi\rho}. \end{aligned}$$

By the way, we have a small reality check here. Moments of inertia should never be negative. □

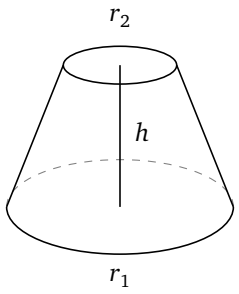

Figure 3.4: A conical frustum with radii $r_1 > r_2$ and height h

Example 3.9. Find the volume of a (circular) **conical frustum** with radii $r_1 > r_2 > 0$ and height h. See Figure 3.4.

Solution. For a lot of regions, we first express the region in rectangular coordinates and then translate the conditions to our new coordinate system. This time let's try to go directly to the cylindrical coordinates for the frustum. Clearly there is no restriction on the angle θ, so we let θ vary from 0 to 2π. Next, it's easy to see that the upper bound for r depends linearly on z. When $z = 0$, r varies from 0 to r_1, and when $z = h$, r varies from 0 to r_2. We can then interpolate linearly between these two ends. Concretely, the side of the frustum is determined by

$$r - r_1 = \frac{r_2 - r_1}{h - 0}(z - 0) \quad \leftrightarrow \quad r = r_1 + \frac{r_2 - r_1}{h} z,$$

so we let r vary from 0 to $r_1 + (r_2 - r_1)z/h$. Finally, z varies from 0 to h, so our frustum F can be written

$$F = \left\{ (r\cos\theta, r\sin\theta, z) \quad : \quad \begin{array}{c} 0 \le \theta \le 2\pi \\ 0 \le z \le h \\ 0 \le r_1 + (r_2 - r_1)z/h \end{array} \right\}.$$

So we have

$$\begin{aligned}
\text{vol}(F) &= \int_F 1 \, dV = \int_0^{2\pi} \int_0^h \int_0^{r_1 + (r_2 - r_1)z/h} r \, dr \, dz \, d\theta \\
&= \int_0^{2\pi} \int_0^h \left[\frac{1}{2} r^2 \right]_0^{r_1 + (r_2 - r_1)z/h} dz \, d\theta = \frac{1}{2} \int_0^{2\pi} \int_0^h \left(r_1 + \frac{r_2 - r_1}{h} z \right)^2 dz \, d\theta \\
&= \frac{1}{2} \int_0^{2\pi} \left[\frac{1}{3} \left(r_1 + \frac{r_2 - r_1}{h} z \right)^3 \frac{h}{r_2 - r_1} \right]_0^h d\theta = \frac{1}{6} \int_0^{2\pi} \frac{h}{r_2 - r_1} \left(r_2^3 - r_1^3 \right) d\theta \\
&= \frac{\pi h}{3} \frac{r_2^3 - r_1^3}{r_2 - r_1} = \boxed{\frac{\pi h}{3}(r_1^2 + r_1 r_2 + r_2^2)}.
\end{aligned}$$

□

3.2.2 Spherical coordinates

The next generalization of polar coordinates to $\mathbb{R}^3$ is slightly more complicated. Given a point (x, y, z) in $\mathbb{R}^3$, we can measure its distance from the origin via

$$\rho = \sqrt{x^2 + y^2 + z^2}.$$

Just like r for polar coordinates, this gives us a radius. Unlike polar coordinates, however, we use two angles to determine the direction in which (x, y, z) is pointing. We first have the *azimuth*, which is the angle made by the vector $(x, y, z)^T$ with the z-axis. We denote the azimuth by φ, and by projecting (x, y, z) onto the z-axis, we see that $z = \rho\cos\varphi$. It remains to determine the x- and y-components, so we project (x, y, z) to the xy-plane to get (x, y). Since $z = \rho\cos\varphi$, one can check that the point (x, y) lies a distance $\rho\sin\varphi$ from the origin. If θ is the angle made by (x, y) with the x-axis (as it was before), then we should be able to convince ourselves that

$$x = \rho\sin\varphi\cos\theta \quad \text{and} \quad y = \rho\sin\varphi\sin\theta.$$

Altogether, spherical coordinates are determined via the transformation $\Phi\colon [0,\infty)_\rho \times [0,2\pi]_\theta \times [0,\pi]_\varphi \to \mathbb{R}^3$ defined by

$$\Phi(\rho,\theta,\varphi) = (\rho\sin\varphi\cos\theta, \rho\sin\varphi\sin\theta, \rho\cos\varphi).$$

Example 3.10. Consider holding ρ constant while θ and φ vary. You should be able to convince yourself that $\Phi(\rho,\theta,\varphi)$ will trace out a sphere of radius ρ. If we fix θ, but vary ρ and φ, we trace out a half-plane with one edge on the z-axis. Fixing φ while ρ and θ run wild will produce a cone with vertex at the origin.

In order to apply the change of variables theorem using spherical coordinates, we'll need to compute the Jacobian determinant $\det\Phi_*$. We have

$$\Phi_* = \begin{pmatrix} \sin\varphi\cos\theta & -\rho\sin\varphi\sin\theta & \rho\cos\varphi\cos\theta \\ \sin\varphi\sin\theta & \rho\sin\varphi\cos\theta & \rho\cos\varphi\sin\theta \\ \cos\varphi & 0 & -\rho\sin\varphi \end{pmatrix},$$

so

$$\begin{aligned}\det\Phi_* &= \cos\varphi \begin{vmatrix} -\rho\sin\varphi\sin\theta & \rho\cos\varphi\cos\theta \\ \rho\sin\varphi\cos\theta & \rho\cos\varphi\sin\theta \end{vmatrix} - \rho\sin\varphi \begin{vmatrix} \sin\varphi\cos\theta & -\rho\sin\varphi\sin\theta \\ \sin\varphi\sin\theta & \rho\sin\varphi\cos\theta \end{vmatrix} \\ &= \cos\varphi(-\rho^2\sin\varphi\cos\varphi) - \rho\sin\varphi(\rho\sin^2\varphi) \\ &= -\rho^2\sin\varphi(\cos^2\varphi + \sin^2\varphi) = -\rho^2\sin\varphi.\end{aligned}$$

Notice that for any $0 \le \varphi \le \pi$, $\sin\varphi \ge 0$, so $|\det\Phi_*| = \rho^2\sin\varphi$. Then, according to the change of variables theorem, we have

$$\boxed{\int_{\Phi(R)} f\ dV = \int_R (f\circ\Phi)|\det\Phi_*|\ dV = \int_R (f\circ\Phi)\rho^2\sin\varphi\ dV},$$

for any region R in $[0,\infty)_\rho \times [0,2\pi]_\theta \times [0,\pi]_\varphi$.

Note. Going forward, you are free to apply the change of variables formula for polar coordinates, cylindrical coordinates, and spherical coordinates without deriving $\det\Phi_*$ again each time.

Example 3.11. For any $\alpha \in (-\pi/2, \pi/2)$, let W_α denote the region which is above the cone $z = \tan\alpha\sqrt{x^2+y^2}$ and inside the sphere of radius 1 centered at the origin in $\mathbb{R}^3$. Find the volume of W_α.

Solution. In rectangular coordinates we have

$$W_\alpha = \{(x,y,z) : \tan\alpha\sqrt{x^2+y^2} \le z, x^2+y^2+z^2 \le 1\}.$$

We can rewrite these conditions in spherical coordinates as

$$(\tan\alpha)\rho\sin\varphi \le \rho\cos\varphi \quad \text{and} \quad \rho \le 1.$$

The first condition simplifies to $\tan\alpha \le \cot\varphi$, which we can further rewrite as $\tan\alpha \le \tan(\frac{\pi}{2}-\varphi)$. Because tan increases monotonically on $(-\pi/2,\pi/2)$, this means that we have $\alpha \le \frac{\pi}{2}-\varphi$. So φ will vary from 0 to $\frac{\pi}{2}-\alpha$. Altogether,

$$W_\alpha = \left\{ \Phi(\rho,\theta,\varphi) \quad : \quad \begin{matrix} 0 \le \theta \le 2\pi \\ 0 \le \rho \le 1 \\ 0 \le \varphi \le \frac{\pi}{2}-\alpha \end{matrix} \right\},$$

so

$$\begin{aligned}\mathrm{vol}(W_\alpha) &= \int_{W_\alpha} 1\ dV = \int_0^{2\pi}\int_0^1\int_0^{\frac{\pi}{2}-\alpha} \rho^2\sin\varphi\ d\varphi\ d\rho\ d\theta \\ &= \int_0^{2\pi}\int_0^1 \left[-\rho^2\cos\varphi\right]_0^{\frac{\pi}{2}-\alpha} d\rho\ d\theta = \int_0^{2\pi}\int_0^1 \rho^2\left(1-\cos\left(\frac{\pi}{2}-\alpha\right)\right) d\rho\ d\theta \\ &= (1-\sin(\alpha))\int_0^{2\pi}\left[\frac{1}{3}\rho^3\right]_0^1 d\theta = (1-\sin(\alpha))\int_0^{2\pi}\frac{1}{3}\ d\theta \\ &= \boxed{\frac{2\pi}{3}(1-\sin(\alpha))}.\end{aligned}$$

Notice that $\text{vol}(W_0) = \frac{2\pi}{3}$, which checks out, since W_0 is half of a ball of radius 1. On the other hand, as α tends towards $\pm\pi/2$, the cone $z = \tan\alpha\sqrt{x^2+y^2}$ tends towards the z-axis. When α is headed to $\pi/2$, W_α is getting smaller and smaller, and indeed $\lim_{\alpha\to(\pi/2)^-}\text{vol}(W_\alpha) = 0$. At the other extreme, as α tends towards $-\pi/2$, W_α is filling up more and more of the unit sphere, and we find that $\lim_{\alpha\to(-\pi/2)^+}\text{vol}(W_\alpha) = 4\pi/3$. □

3.2.3 Transformations that go the wrong way

We can recap the change of variables theorem as follows: given a coordinate system $\Phi\colon D \to R$, integrals over R can be "pulled back" via Φ to integrals over D, and we pick up a scale factor of $|\det\Phi_*|$. Sometimes, however, we'd really prefer to *push* our integrals rather than pull them. Specifically, instead of the coordinate system Φ above, we might have its inverse[5] $\Psi\colon R \to D$. We might want to move integrals from R to D, but find Φ to difficult to compute explicitly. For this, the following two facts are helpful:

1. if $\Psi = \Phi^{-1}$, then $\Psi_* = \Phi_*^{-1}$;
2. if $\det\Phi_* \neq 0$, then $(\det\Phi_*)(\det\Phi_*^{-1}) = 1$.

From these facts we see that the stretch factor $|\det\Phi_*|$ is given by $|\det\Psi_*|^{-1}$. This allows us to rewrite the change of variables theorem as

$$\boxed{\int_R f\ dV = \int_D (f\circ\Phi)|(\det\Psi_*)\circ\Phi|^{-1}\ dV = \int_D (f\,|\det\Psi_*|^{-1})\circ\Phi\ dV.} \tag{3.2}$$

Note that in the middle expression we precompose $\det\Psi_*$ with with Φ. This makes more sense if we assign coordinate names. Say we're working in two dimensions and our coordinate systems are

$$\Phi\colon D_{u,v} \to R_{x,y} \quad \text{and} \quad \Psi\colon R_{x,y} \to R_{u,v}.$$

Then $\det\Psi_*$ will be a function of x and y. But since we want to integrate over D, we need a function of u and v — this is what we accomplish by precomposing with Φ.

The upshot is that, in sufficiently nice circumstances[6], we can pull our integral back over Φ without ever having an explicit formula for Φ. Let's see this in action.

Example 3.12. Evaluate the integral

$$\int_R xy\ dA,$$

where R is the region in the first quadrant that is enclosed by the hyperbolas $x^2-y^2=1$, $x^2-y^2=4$, and the circles $x^2+y^2=9$, $x^2+y^2=16$.

Solution. The bounds for R are very suggestive of a transformation $\Psi\colon R_{x,y} \to D_{u,v}$:

$$\Psi(x,y) := (x^2-y^2, x^2+y^2).$$

This allows us to write R as the image of a rectangle:

$$R = \{\Psi^{-1}(u,v) : 1 \le u \le 4, 9 \le v \le 16\}.$$

In order to directly apply the change of variables theorem at this point, we would need to come up with an expression for $\Phi = \Psi^{-1}$. Instead, let's try to apply Equation 3.2. We have

$$\Psi_* = \begin{pmatrix} 2x & -2y \\ 2x & 2y \end{pmatrix},$$

so $\det\Psi_* = 8xy$. Now Equation 3.2 asks us to think about $f\circ\Phi$ and $(\det\Psi_*)\circ\Phi$. Since we don't know Φ, this can be challenging. But the rightmost expression of Equation 3.2 explains how this will typically work out: we first think about the product $f\,|\det\Psi_*|$ in the coordinates of R, and then precompose with Φ. In this example we have

$$f(x,y)\cdot(\det\Psi_*(x,y))^{-1} = (xy)(8xy)^{-1} = \frac{1}{8}.$$

[5]Recall that coordinate systems are required to be injective on the interior of their domain, so, assuming that $\Phi(\mathring{D}) = \mathring{R}$, we can invert Φ on $\mathring{R}$.

[6]Since we don't have our hands on Φ, the nice circumstance is usually that the expression $f\,|\det\Psi_*|^{-1}$ quickly simplifies to something that's easy to recognize in the coordinates of D.

Of course precomposing the constant $\frac{1}{8}$ with Φ will do nothing, so we have

$$\int_R xy\ dA = \int_9^{16}\int_1^4 \frac{1}{8}\ du\ du v = 7\cdot 16\cdot\frac{1}{8} = \boxed{14}.$$

□

Finally, here's an example of the same phenomenon in $\mathbb{R}^3$.

Example 3.13. Find the volume of the oblique parallelepiped that is bounded by the planes

$$x+y+2z=\pm 3,\quad x-2y+z=\pm 2,\quad \text{and}\quad 4x+y+z=\pm 6.$$

Solution. Once again, our region suggest a transformation $\Psi\colon R_{x,y,z}\to D_{u,v,w}$:

$$\Psi(x,y,z) := (x+y+2z, x-2y+z, 4x+y+z).$$

This time, Ψ is a linear transformation, so finding an explicit formula for Φ isn't that difficult. But we like to be lazy, so let's use Equation 3.2. Computing Ψ_* is easy:

$$\Psi_* = \begin{pmatrix} 1 & 1 & 2\\ 1 & -2 & 1\\ 4 & 1 & 1\end{pmatrix}.$$

Then

$$\begin{aligned}\det\Psi_* &= \begin{vmatrix}-2 & 1\\ 1 & 1\end{vmatrix} - \begin{vmatrix}1 & 2\\ 1 & 1\end{vmatrix} + 4\begin{vmatrix}1 & 2\\ -2 & 1\end{vmatrix}\\ &= -3-(-1)+4(5) = 18.\end{aligned}$$

So we find that

$$\begin{aligned}\operatorname{vol}(R) &= \int_R 1\ dV = \int_D |(\det\Psi_*)\circ\Phi|^{-1}\ dV\\ &= \int_{-3}^{3}\int_{-2}^{2}\int_{-6}^{6}\frac{1}{18}\ dw\ dv\ du = (6)(4)(12)\frac{1}{18} = \boxed{16},\end{aligned}$$

where we use the fact that $D = [-3,3]_u\times[-2,2]_v\times[-6,6]_w\subset\mathbb{R}^3$. □

4 Week 4

This week begins with a midterm, after which we'll begin turning our attention to vector calculus. We define vector fields and recall a vector-valued derivative of scalar-valued functions. Afterwards, we define two new derivatives involving vector fields and consider the problem of antidifferentiation for these derivatives. Finally, we give a very brief hint at how antidifferentiation can detect nontrivial *topology* in a subset of $\mathbb{R}^n$.

4.1 Day 9: Midterm 1

Goals

By the end of today's class, we should be able to do the following.

1. Solve the problems on the midterm.

Here are the midterm 1 problems and solutions.

1. Prove that Riemann integration over compact boxes is preserved by linear combination. That is, let $B \subset \mathbb{R}^n$ be a compact box, let c and d be real numbers, and let $f, g : B \to \mathbb{R}$ be functions which are integrable over B. Prove that the function $h : B \to \mathbb{R}$ defined by $h(x) := c \cdot f(x) + d \cdot g(x)$ is integrable over B.

Solution. We use the integrability criterion to prove that h is integrable. To this end, choose $\epsilon > 0$. Our goal is to construct a partition $\mathcal{P}$ of B for which we have

$$U(h, \mathcal{P}) - L(h, \mathcal{P}) < \epsilon. \tag{4.1}$$

If $c = 0$ and $d = 0$, this inequality obviously holds, since all lower and upper sums of h are 0. So we assume that c and d are not both 0. Let us first consider the case where $d = 0$. Then we define $\epsilon_f := \epsilon/|c|$. Because f is integrable, the integrability criterion tells us that there is a partition $\mathcal{P}$ of B such that

$$U(f, \mathcal{P}) - L(f, \mathcal{P}) < \epsilon_f.$$

That is,

$$\sum_J (M_J(f) - m_J(f))\text{vol}(J) < \frac{\epsilon}{|c|}.$$

Multiplying both sides of this inequality by $|c|$, we find that

$$\begin{aligned} \epsilon > |c| \sum_J (M_J(f) - m_J(f))\text{vol}(J) &= \sum_J (M_J(|c|f) - m_J(|c|f))\text{vol}(J) \\ = \sum_J (M_J(c \cdot f) - m_J(c \cdot f))\text{vol}(J) &= U(h, \mathcal{P}) - L(h, \mathcal{P}). \end{aligned}$$

That is, (4.1) holds, so h is integrable. Notice that the penultimate equality holds because, if $c < 0$, then

$$\begin{aligned} M_J(c \cdot f) - m_J(c \cdot f) &= (-m_J(-c \cdot f)) - (-M_J(-c \cdot f)) \\ &= -m_J(|c|f) + M_J(|c|f) \\ &= M_J(|c|f) - m_J(|c|f), \end{aligned}$$

for any subbox J associated to $\mathcal{P}$. The fact that h is integrable when $c = 0$ and $d \neq 0$ follows by a completely analogous argument.

Our primary concern is the case where neither c nor d is zero. In this case we define

$$\epsilon_f := \frac{\epsilon}{2|c|} \quad \text{and} \quad \epsilon_g := \frac{\epsilon}{2|d|}.$$

Because both f and g are integrable, the integrability criterion allows us to choose partitions $\mathcal{P}_f$ and $\mathcal{P}_g$ of B such that

$$U(f,\mathcal{P}_f)-L(f,\mathcal{P}_f)<\epsilon_f \quad \text{and} \quad U(g,\mathcal{P}_g)-L(g,\mathcal{P}_g)<\epsilon_g.$$

Now we can let $\mathcal{P}$ be the common refinement of $\mathcal{P}_f$ and $\mathcal{P}_g$. For each subbox J associated to $\mathcal{P}$, we have

$$\begin{aligned} M_J(h)-m_J(h) &= M_J(c\cdot f+d\cdot g)-m_J(c\cdot f+d\cdot g)\\ &\le (M_J(c\cdot f)-m_J(c\cdot f))+(M_J(d\cdot g)-m_J(d\cdot g)), \end{aligned}$$

loosely because the largest value attained by a sum of two functions is at most the sum of the two functions' largest values, and similarly for smallest values. (The analogous property holds for suprema and infima.) So we have

$$\begin{aligned} M_J(h)-m_J(h) &\le (M_J(|c|\cdot f)-m_J(|c|\cdot f))+(M_J(|d|\cdot g)-m_J(|d|\cdot g))\\ &= |c|(M_J(f)-m_J(f))+|d|(M_J(g)-m_J(g)) \end{aligned} \tag{4.2}$$

because, as explained above, $M_J(c\cdot f)-m_J(c\cdot f)=M_J(|c|f)-m_J(c|f|)$, for any $c\in\mathbb{R}$, and similarly for d. Now we multiply (4.2) by $\mathrm{vol}(J)$ and sum over all subboxes J associated to $\mathcal{P}$ to see that

$$\sum_J (M_J(h)-m_J(h))\mathrm{vol}(J) \le |c|\sum_J (M_J(f)-m_J(f))\mathrm{vol}(J)+|d|\sum_J (M_J(g)-m_J(g))\mathrm{vol}(J).$$

But this just means that

$$U(h,\mathcal{P})-L(h,\mathcal{P}) \le |c|(U(f,\mathcal{P})-L(f,\mathcal{P}))+|d|(U(g,\mathcal{P})-L(g,\mathcal{P})).$$

Finally, $\mathcal{P}$ is a refinement of $\mathcal{P}_f$. Lemma 1.9 of the class notes tells us that $L(f,\mathcal{P}_f)\le L(f,\mathcal{P})$ and $U(f,\mathcal{P})\le U(f,\mathcal{P}_f)$, so

$$U(f,\mathcal{P})-L(f,\mathcal{P}) \le U(f,\mathcal{P}_f)-L(f,\mathcal{P}_f).$$

Similarly,

$$U(g,\mathcal{P})-L(g,\mathcal{P}) \le U(g,\mathcal{P}_g)-L(g,\mathcal{P}_g).$$

So we find that

$$\begin{aligned} U(h,\mathcal{P})-L(h,\mathcal{P}) &\le |c|(U(f,\mathcal{P})-L(f,\mathcal{P}))+|d|(U(g,\mathcal{P})-L(g,\mathcal{P}))\\ &\le |c|(U(f,\mathcal{P}_f)-L(f,\mathcal{P}_f))+|d|(U(g,\mathcal{P}_g)-L(g,\mathcal{P}_g))\\ &< |c|\cdot\epsilon_f+|d|\cdot\epsilon_g\\ &= |c|\frac{\epsilon}{2|c|}+|d|\frac{\epsilon}{2|d|}=\epsilon, \end{aligned}$$

meaning that (4.1) is satisfied. According to the integrability criterion, h is integrable.

Note. An even better way to do solve the problem is to (a) show that cf is integrable, for any $c\in\mathbb{R}$, (b) show that $f+g$ is integrable, and (c) combine these two facts to get the desired conclusion. □

2. Let B be a compact box in $\mathbb{R}^n$, let $f:B\to\mathbb{R}$ be a function on B, and suppose that for each $n\in\mathbb{N}$ we have a function $f_n:B\to\mathbb{R}$. We say that the sequence (f_n) *converges uniformly to* f if, for every $\epsilon>0$, there is some $N\in\mathbb{N}$ such that

$$\sup\{|f(x)-f_n(x)| : x\in B\}<\epsilon,$$

whenever $n\ge N$. That is, the difference between $f(x)$ and $f_n(x)$ is at most ϵ, for all $x\in B$. Suppose that (f_n) converges uniformly to f, and that each f_n is integrable over B. Prove that f is also integrable over B.

Solution. We will prove that f is integrable using the integrability criterion. To this end, choose $\epsilon>0$. With $\epsilon>0$ chosen, define

$$\epsilon' := \frac{\epsilon}{4\mathrm{vol}(B)} \quad \text{and} \quad \epsilon'' := \frac{\epsilon}{2}.$$

Because (f_n) converges uniformly to f, we can choose $N \in \mathbb{N}$ so that

$$\sup\{|f(x) - f_n(x)| : x \in B\} < \epsilon',$$

for all $n \geq N$. Moreover, since f_N is integrable over B, the integrability criterion tells us that we can choose a partition $\mathcal{P}$ of B so that

$$U(f_N, \mathcal{P}) - L(f_N, \mathcal{P}) < \epsilon''.$$

Now let's think about the difference $U(f, \mathcal{P}) - L(f, \mathcal{P})$. Since $|f(x) - f_N(x)| < \epsilon'$, we know that $f_N(x) - \epsilon' < f(x) < f_N(x) + \epsilon'$, for all $x \in B$. In particular, this means that for any subbox J of B associated to $\mathcal{P}$, we have

$$M_J(f) \leq M_J(f_N) + \epsilon' \quad \text{and} \quad m_J(f) \geq m_J(f_N) - \epsilon'.$$

It follows that

$$M_J(f) - m_J(f) \leq (M_J(f_N) + \epsilon') - (m_J(f_N) - \epsilon') = (M_J(f_N) - m_J(f_N)) + 2\epsilon',$$

for every subbox J associated to $\mathcal{P}$. We can multiply this inequality by $\text{vol}(J)$ and sum over all subboxes associated to $\mathcal{P}$:

$$\sum_J (M_J(f) - m_J(f))\text{vol}(J) \leq \sum_J (M_J(f_N) - m_J(f_N))\text{vol}(J) + 2\epsilon' \sum_J \text{vol}(J).$$

That is,

$$\begin{aligned} U(f, \mathcal{P}) - L(f, \mathcal{P}) &\leq U(f_N, \mathcal{P}) - L(f_N, \mathcal{P}) + 2\epsilon'\text{vol}(B) \\ &< \epsilon'' + \frac{\epsilon}{2} = \epsilon. \end{aligned}$$

So for any $\epsilon > 0$, we can find a partition $\mathcal{P}$ of B so that the difference between the upper and lower sums of f over $\mathcal{P}$ is less than ϵ. The integrability criterion allows us to conclude that f is integrable over $\mathcal{P}$. □

3. For any $n \geq 0$, the **standard n-simplex** is the region

$$\Delta^n := \left\{ (t_1, \ldots, t_n) \in \mathbb{R}^n : \sum_{i=1}^n t_i \leq 1 \text{ and } t_1, \ldots, t_n \geq 0 \right\}.$$

Use multiple integration[1] to prove that the barycenter (or centroid) of Δ^4 is $(\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5})$.

Solution. Remember that the barycenter is given by $\overline{t} = (\overline{t}_1, \overline{t}_2, \overline{t}_3, \overline{t}_4)$, where

$$\overline{t}_i = \frac{1}{\text{vol}(\Delta^4)} \int_{\Delta^4} t_i \, dV = 24 \int_{\Delta^4} t_i \, dV.$$

for $1 \leq i \leq 4$. Now to apply Fubini's theorem, we need to rewrite the inequalities defining our region Δ^4. Let's say we want t_1 to be the variable of our outermost integration. Then we notice that t_1 can be as small as 0 or as large as 1 and have $0 \leq t_1 \leq 1$. Next, suppose we have chosen a value t_1 and want to know what values t_2 is allowed to attain. Certainly $t_2 \geq 0$, and the condition $t_1 + t_2 + t_3 + t_4 \leq 1$ tells us that $t_2 \leq 1 - t_1$, since we know that neither t_3 nor t_4 is negative. So

$$0 \leq t_1 \leq 1 \quad \text{and} \quad 0 \leq t_2 \leq 1 - t_2.$$

Similarly, once we know t_1 and t_2, the largest value t_3 is allowed to achieve is $1 - t_1 - t_2$. Finally, t_4 runs from 0 to $1 - t_1 - t_2 - t_3$. Altogether,

$$\overline{t}_i = 24 \int_0^1 \int_0^{1-t_1} \int_0^{1-t_1-t_2} \int_0^{1-t_1-t_2-t_3} t_i \, dt_4 \, dt_3 \, dt_2 \, dt_1.$$

[1] There are certainly other (perhaps better) ways to find this barycenter, but you are required to use multiple integration.

Here the 24 at the front comes from the fact that $\frac{1}{\text{vol}(\Delta^4)} = \frac{1}{(1/4!)} = 24$. Then we can compute $\overline{t}_4$:

$$\begin{aligned}
\overline{t}_4 &= 24\int_0^1\int_0^{1-t_1}\int_0^{1-t_1-t_2}\int_0^{1-t_1-t_2-t_3} t_4\,dt_4\,dt_3\,dt_2\,dt_1 \\
&= 24\int_0^1\int_0^{1-t_1}\int_0^{1-t_1-t_2}\left[\frac{1}{2}t_4^2\right]_0^{1-t_1-t_2-t_3} dt_3\,dt_2\,dt_1 \\
&= 12\int_0^1\int_0^{1-t_1}\int_0^{1-t_1-t_2}(1-t_1-t_2-t_3)^2\,dt_3\,dt_2\,dt_1 \\
&= 12\int_0^1\int_0^{1-t_1}\left[-\frac{1}{3}(1-t_1-t_2-t_3)^3\right]_0^{1-t_1-t_2} dt_2\,dt_1 \\
&= 4\int_0^1\int_0^{1-t_1}(1-t_1-t_2)^3\,dt_2\,dt_1 \\
&= \int_0^1\left[-(1-t_1-t_2)^4\right]_0^{1-t_1} dt_1 \\
&= \int_0^1(1-t_1)^4\,dt_1 \\
&= \left[-\frac{1}{5}(1-t_1)^5\right]_0^1 = \frac{1}{5}.
\end{aligned}$$

Yikes. That was a lot of work to get one coordinate. We obtain the other coordinates via the symmetry of Δ^4. Namely, consider the C^1-transformation $\Phi\colon \mathbb{R}^4 \to \mathbb{R}^4$ defined by

$$\Phi(t_1,t_2,t_3,t_4) := (t_1,t_2,t_4,t_3).$$

Then

$$\Phi_* = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix},$$

so $|\det\Phi_*| \equiv 1$. Clearly Φ is a bijection between the boxes $\mathbb{R}^4$ and $\mathbb{R}^4$, and since $|\det\Phi_*|$ is nowhere-vanishing, we see that Φ is a coordinate system. Moreover, one can easily compute that $\Phi(\Delta^4) = \Delta^4$, and thus

$$\begin{aligned}
\overline{t}_3 &= \int_{\Delta^4} t_3\,dV = \int_{\Phi^{-1}(\Delta^4)} (t_3\circ\Phi)|\det\Phi_*|\,dV \\
&= \int_{\Delta^4} t_4\,dV = \overline{t}_4 = \frac{1}{5}.
\end{aligned}$$

By completely analogous reasoning, the symmetry of Δ^4 ensures that $\overline{t}_1 = \overline{t}_2 = \overline{t}_4 = \frac{1}{5}$. So we conclude that our barycenter is given by $\overline{t} = (1/5, 1/5, 1/5, 1/5)$.

Remark. We could probably also solve this problem by putting a coordinate system on Δ^4 whose domain is a (hyper)cube, and then using the change of variables theorem. This should lead to the same result, but I'm too lazy to write it up right now. □

4. For $a \in \mathbb{R}$, consider the region

$$W_a = \{(x,y,z) : (x-a)^2 + y^2 \le 1, (x-a)^2 + y^2 - 1 \le z \le 1 - (x-a)^2 - y^2\}.$$

Prove that the moment of inertia of W_a about the z-axis is given by $\pi a^2 + \pi/3$.

Solution. Notice that W_a is symmetric about the axis $x = a$. Based on this, we write down a shifted version of cylindrical coordinates:

$$\Phi(r,\theta,z) := (r\cos\theta + a, r\sin\theta, z).$$

Notice that r isn't our usual radius: it's the radius from the axis $x = a$, and it satisfies

$$r = \sqrt{(x-a)^2 + y^2}.$$

But this is nice, because it means we can rewrite the first inequality of W_a as $0 \le r \le 1$ and write the second inequality as $r^2 - 1 \le z \le 1 - r^2$. So

$$W_a = \{\Phi(r,\theta,z) : 0 \le r \le 1, r^2 - 1 \le z \le 1 - r^2\},$$

making integrals over W_a pretty straightforward to compute. Namely, we have

$$\Phi_* = \begin{pmatrix} \cos\theta & -r\sin\theta & 0 \\ \sin\theta & r\cos\theta & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

so

$$\det \Phi_* = \begin{vmatrix} \cos\theta & -r\sin\theta \\ \sin\theta & r\cos\theta \end{vmatrix} = r\cos^2\theta + r\sin^2\theta = r.$$

We conclude that $\Phi : [0,\infty)_r \times [0,2\pi]_\theta \times \mathbb{R}_z \to \mathbb{R}^3_{x,y,z}$ is a coordinate system. Indeed:

- the domain $[0,\infty) \times [0,2\pi] \times \mathbb{R}$ is a box;
- on the interior $(0,\infty) \times (0,2\pi) \times \mathbb{R}$, Φ is injective;
- the image of Φ is all of $\mathbb{R}^3_{x,y,z}$;
- the Jacobian determinant $\det \Phi_* = r$ does not vanish on the interior of the domain of Φ.

This allows us to compute the moment of inertia of W_a about the z-axis:

$$I_z = \int_{W_a} D^2 \, dV = \int_0^{2\pi} \int_0^1 \int_{r^2-1}^{1-r^2} D^2(\Phi(r,\theta,z)) \, r \, dz \, dr \, d\theta.$$

The only catch is that we need to write down $D^2(\Phi(r,\theta,z))$. Certainly $D^2(x,y,z) = x^2 + y^2$, so we have

$$D^2(\Phi(r,\theta,z)) = (r\cos\theta + a)^2 + r^2\sin^2\theta = r^2 + 2ar\cos\theta + a^2.$$

Finally, we can integrate:

$$\begin{aligned}
\int_{W_a} D^2 \, dV &= \int_0^1 \int_0^{2\pi} \int_{r^2-1}^{1-r^2} (r^2 + 2ar\cos\theta + a^2) \, r \, dz \, d\theta \, dr \\
&= \int_0^1 \int_0^{2\pi} (r^2 + 2ar\cos\theta + a^2) \, r \, (2 - 2r^2) \, d\theta \, dr \\
&= 2\pi \int_0^1 (r^2 + a^2) \, r \, (2 - 2r^2) \, dr \\
&= 2\pi \int_0^1 (2a^2 r + 2r^3 - 2a^2 r^3 - 2r^5) \, dr \\
&= 2\pi \left[a^2 r^2 + \frac{1}{2} r^4 - \frac{1}{2} a^2 r^4 - \frac{1}{3} r^6 \right]_0^1 \\
&= 2\pi \left(\frac{a^2}{2} + \frac{1}{6} \right) = \boxed{\pi a^2 + \frac{\pi}{3}}.
\end{aligned}$$

Notice that along the way we swapped the order of integration, and used the identity $\int_0^{2\pi} \cos\theta \, d\theta = 0$. We have a reality check here: the minimum moment of inertia occurs when $a = 0$, so that the cylinder is centered on the z-axis, and increases quadratically as we move the cylinder away from the axis.

An alternative argument uses the *parallel axis theorem*, established in Homework 2. This says that if L_0 is an axis passing through the center of gravity of W_a and L is a parallel axis at a distance a from L_0, then

$$I = I_0 + Ma^2,$$

where I and I_0 are the moments of inertia of W_a about L and L_0, respectively, and M is the mass of W_a. To use this approach, we must

(a) compute the mass M of W_a;
(b) determine the center of gravity of W_a;
(c) determine the moment of inertia I_0 of W_a about the axis which passes through the center of gravity and is parallel to the z-axis.

I won't fill out the details here, but this is a perfectly viable strategy. □

Whew! With that behind us, we can say goodbye to ϵ for a while.

4.2 Day 10: Gradients and curls

Goals

By the end of today's class, we should be able to do the following.

1. Use the chain rule to define the **gradient** of a function of several variables.
2. Explain some properties — mathematically and physically — which distinguish gradient vector fields within the larger class of vector fields on $\mathbb{R}^n$.
3. Define the **curl** of a vector field, and use this to state a necessary condition for a C^1-vector field on $\mathbb{R}^n$ to be a gradient vector field.

So far in this course we've focused on integration — a good thing to do in a course on integral calculus. But the Fundamental Theorem of Calculus tells us that integration is very closely related to antidifferentiation; a major theme of the next few weeks will be to find antiderivatives for vector fields. To do this, we'll first need to define **vector fields**, and then we'll define some derivatives which output vector fields (so that we can then *anti*differentiate).

Definition. Given a subset A of $\mathbb{R}^n$, a **vector field** X on A is a vector-valued function $X : A \to \mathbb{R}^n$. We say that X is **continuous** if each component function is continuous, and we say that X is a **$\mathbf{C^1}$-vector field** if each component function is continuously differentiable.

Remark. Unless otherwise stated, we'll assume that the vector fields we consider are C^1.

We will typically use component notation to write our vector fields. That is, we will write

$$X = \langle F_1, \dots, F_n \rangle,$$

where $F_i : A \to \mathbb{R}$ is a scalar-valued function of several variables. For vector fields on subsets of $\mathbb{R}^2$ and $\mathbb{R}^3$ we may also use the notation

$$X = F_1\mathbf{i} + F_2\mathbf{j} \quad \text{or} \quad X = F_1\mathbf{i} + F_2\mathbf{j} + F_3\mathbf{k},$$

where $\mathbf{i}$, $\mathbf{j}$, and $\mathbf{k}$ are the constant vector fields given in components by

$$\mathbf{i} = \langle 1, 0, 0 \rangle, \quad \mathbf{j} = \langle 0, 1, 0 \rangle, \quad \mathbf{k} = \langle 0, 0, 1 \rangle.$$

Of course, when treating $\mathbf{i}$ and $\mathbf{j}$ as vector fields on $\mathbb{R}^2$ we omit the third component.

We're already familiar with one derivative that produces a vector field: the gradient. Let's state the definition again, though, because my preferred definition is maybe a little different from what you saw in 32A.

Definition. The **gradient** of a differentiable function $f : \mathbb{R}^n \to \mathbb{R}$ is the unique vector field $\nabla f : \mathbb{R}^n \to \mathbb{R}^n$

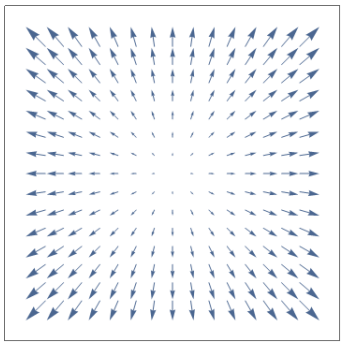

(a) $X = \langle x, y \rangle$

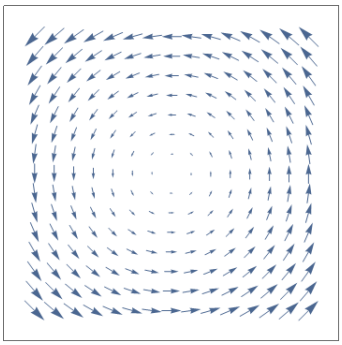

(b) $Y = \langle -y, x \rangle$

Figure 4.1: A pair of vector fields on $\mathbb{R}^2$

which satisfies the differential equation

$$\frac{d}{dt}\left(f(\mathbf{r}(t))\right) = \nabla f(\mathbf{r}(t)) \cdot \mathbf{r}'(t),$$

for all vector-valued functions $\mathbf{r}\colon \mathbb{R} \to \mathbb{R}^n$ and all $t \in \mathbb{R}$.

Exercise 4.1. Use the chain rule to prove that, in rectangular coordinates, the gradient of $f : \mathbb{R}^n \to \mathbb{R}$ is given by

$$\nabla f = \sum_{k=1}^{n} \frac{\partial f}{\partial x^k} \frac{\partial}{\partial x^k},$$

where $\frac{\partial}{\partial x^k}$ denotes the constant vector field $\mathbf{e}_k$ on $\mathbb{R}^n$. In the (optional portion of the) homework you'll see that this sort of formula does not necessarily hold in non-rectangular coordinate systems.

Since the gradient is a derivative for scalar-valued functions which produces a vector field, we can then ask for the *antiderivative* of a vector field. Inspired by physics, we call this antiderivative the *scalar potential*.

Definition. Let $X : A \to \mathbb{R}^n$ be a vector field on a subset A of $\mathbb{R}^n$. A **scalar potential** (or **scalar potential field**) is a function $\phi : A \to \mathbb{R}$ satisfying

$$\nabla\phi(p) = X(p),$$

for all points $p \in A$. If a scalar potential exists for X, we say that X is **conservative**.

Now given an arbitrary C^1-vector field X, there is no reason to believe that X should be conservative. But conservative vector fields are particularly special. Over the next couple of weeks, we will discuss several distinguishing properties of conservative vector fields (including why they're called *conservative*), but for now we point out that conservative vector fields point in the direction of greatest increase for their scalar potentials. That is, *conservative vector fields attempt to maximize some objective function*. So, if our minds are not shrouded by actual knowledge of physics, we can imagine a physicist watching some objects[2] move around, following the flow of a vector field X. The physicist might then wonder to herself, "Are these objects trying to maximize some function?" This isn't *so* far-fetched: most objects are trying to minimize their potential energy, so the physicist is wondering, "What's the potential energy function in my setting?" In mathematical terms, she's asking, "Does the vector field these objects are following admit a scalar potential?" So she's looking for an antiderivative, and as mathematicians we're thrilled that someone cares about our work.

Hopefully we're now convinced that this antidifferentiation problem is worth solving; next, let's try to develop some tools to solve it. Consider the vector fields

$$X = \langle x, y \rangle \quad \text{and} \quad Y = \langle -y, x \rangle$$

on $\mathbb{R}^2$, depicted in Figure 4.1. Based on the plots, it seems as though X is trying to maximize a quantity — namely, distance from the origin — and therefore could be a gradient. On the other hand, Y has *closed orbits*.

[2] Let's take full advantage of our ignorance of physics and say she's watching "particles."

That is, we can write down a parametrized curve $\mathbf{r}\colon [0,1] \to \mathbb{R}^2$ such that

$$Y(\mathbf{r}(t)) = \mathbf{r}'(t) \quad \text{and} \quad \mathbf{r}(0) = \mathbf{r}(1).$$

For instance, $\mathbf{r}(t) = (\cos(2\pi t), \sin(2\pi t))$ will work. If $Y = \nabla\phi$, then we would find that

$$\frac{d}{dt}(\phi(\mathbf{r}(t))) = \nabla\phi(\mathbf{r}(t)) \cdot \mathbf{r}'(t) = \|Y(\mathbf{r}(t))\|^2,$$

for all t. For the particular closed curve we've chosen, $Y(\mathbf{r}(t))$ never vanishes, so this derivative is always positive. But this is nonsense: we know that $\mathbf{r}(0) = \mathbf{r}(1)$, so $\phi(\mathbf{r}(0)) = \phi(\mathbf{r}(1))$. By assumption, ϕ is differentiable, so the mean value theorem tells us that there must be some value $t_0 \in (0,1)$ where $\frac{d}{dt}(\phi(\mathbf{r}(t_0))) = 0$. If Y were a gradient, then $(\mathbf{r}(t), \phi(\mathbf{r}(t)))$ would parametrize some kind of *Relativity* staircase.

With this vague discussion behind us, it's really not difficult to prove that Y fails to be a gradient. If $Y = \nabla\phi$, then

$$-y = \phi_x \quad \text{and} \quad x = \phi_y.$$

From this it follows that $\phi_{xy} = -1$, while $\phi_{yx} = 1$. Of course this violates Clairaut's theorem[3], and we conclude that no such ϕ exists. On the other hand, if $X = \nabla\phi$, then

$$x = \phi_x \quad \text{and} \quad y = \phi_y.$$

This checks out, as it tells us that $\phi_{xy} = 0 = \phi_{yx}$. So we have a necessary condition for a vector field to be conservative.

Proposition 4.2. *Let $X : A \to \mathbb{R}^2$ be a C^1-vector field on an open subset A of $\mathbb{R}^2$, and write $X = \langle F_1, F_2 \rangle$. If X is conservative, then*

$$\frac{\partial F_1}{\partial y} = \frac{\partial F_2}{\partial x}, \tag{4.3}$$

for all points $(x,y) \in A$.

A word on necessary and sufficient conditions. In our search for a characterization of conservative vector fields, Proposition 4.2 give us what is known as a **necessary condition**. That is, in order for X to be conservative, its component functions must satisfy (4.3). On the other hand, we sometimes look for **sufficient conditions**. These are properties which, when satisfied by X, guarantee that X is a conservative vector field. In the best of all worlds, we would find a condition which is both necessary and sufficient, and thus provides a perfect test of conservativity.

Exercise 4.3. Prove that the vector field

$$X = \left\langle \frac{-y}{x^2+y^2}, \frac{x}{x^2+y^2} \right\rangle,$$

defined on the open subset $\mathbb{R}^2 \setminus \{(0,0)\}$ of $\mathbb{R}^2$, satisfies Equation 4.3, but is not conservative. This demonstrates that the necessary condition given by Proposition 4.2 is not sufficient.

At present, the only sufficient condition we have for conservative vector fields is the definition itself. So the only way to verify that a vector field is conservative is to explicitly produce a scalar potential. For our example of $X = \langle x, y \rangle$, this is not so bad. As we said above, any ϕ satisfying $X = \nabla\phi$ has

$$x = \phi_x \quad \text{and} \quad y = \phi_y.$$

From the first equation we deduce that

$$\phi(x,y) = \frac{1}{2}x^2 + g(y),$$

for some function $g : \mathbb{R} \to \mathbb{R}$. Notice that where in single-variable calculus we would add an arbitrary constant to an antiderivative, here we must add an arbitrary function of the remaining variable (in this case y). This is

[3] Recall that Clairaut's theorem tells us that, if all second-order derivatives of $\phi : \mathbb{R}^n \to \mathbb{R}$ exist and are continuous, then $\phi_{x^i x^j} = \phi_{x^j x^i}$, for all $1 \le i, j \le n$. This is also called *equality of mixed partials*.

because the x-derivative ϕ_x doesn't see such functions. Differentiating this expression for ϕ against y, we find that

$$y = \phi_y = \frac{\partial}{\partial y}\left(\frac{1}{2}x^2 + g(y)\right) = g'(y),$$

so $g(y) = \frac{1}{2}y^2 + C$, for some constant C. That is,

$$\phi(x,y) = \frac{1}{2}x^2 + \frac{1}{2}y^2 + C,$$

for some arbitrary constant C. So $X = \langle x, y\rangle$ is indeed conservative, and we see that flowing along X will increase (one half of the square of) distance from the origin.

Proposition 4.2 gives us a necessary condition for conservative vector fields on open subsets of $\mathbb{R}^2$. Using the same reasoning, we can develop a necessary condition for conservative vector fields on $\mathbb{R}^3$. Let's consider a vector field

$$X = \langle F_1, F_2, F_3\rangle,$$

and suppose $X = \nabla\phi$, for some scalar potential $\phi : \mathbb{R}^3 \to \mathbb{R}$. Then

$$\langle F_1, F_2, F_3\rangle = \left\langle \frac{\partial\phi}{\partial x}, \frac{\partial\phi}{\partial y}, \frac{\partial\phi}{\partial z}\right\rangle.$$

Once again, we can use Clairaut's theorem to make some deductions about the component functions F_1, F_2, and F_3. But this time we get three differential equations rather than one:

$$\begin{aligned}
\frac{\partial F_3}{\partial y} &= \frac{\partial^2\phi}{\partial y\partial z} = \frac{\partial^2\phi}{\partial z\partial y} = \frac{\partial F_2}{\partial z}\\
\frac{\partial F_1}{\partial z} &= \frac{\partial^2\phi}{\partial z\partial x} = \frac{\partial^2\phi}{\partial x\partial z} = \frac{\partial F_3}{\partial x}\\
\frac{\partial F_2}{\partial x} &= \frac{\partial^2\phi}{\partial x\partial y} = \frac{\partial^2\phi}{\partial y\partial x} = \frac{\partial F_1}{\partial y}.
\end{aligned}$$

Whoa. We can repackage these three differential equations as a single vector equation:

$$\left\langle \frac{\partial F_3}{\partial y} - \frac{\partial F_2}{\partial z}, \frac{\partial F_1}{\partial z} - \frac{\partial F_3}{\partial x}, \frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y}\right\rangle = \mathbf{0}. \tag{4.4}$$

The order in which we listed the three differential equations may seem mysterious, but notice that, if we're willing to abuse notation a bit, we can write the vector field on the left hand side of Equation 4.4 as a cross product $\nabla \times X$, where

$$\nabla = \left\langle \frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z}\right\rangle.$$

This vector field will be especially important to us in the coming weeks, so let's give it a name.

Definition. Let $X : A \to \mathbb{R}^3$ be a C^1-vector field on some open subset A of $\mathbb{R}^3$. The **curl** of X, denoted $\mathrm{curl}X$ or $\nabla \times X$, is a vector field on A defined by

$$\mathrm{curl}X := \nabla \times X := \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \partial/\partial x & \partial/\partial y & \partial/\partial z \\ F_1 & F_2 & F_3 \end{vmatrix} = \left\langle \frac{\partial F_3}{\partial y} - \frac{\partial F_2}{\partial z}, \frac{\partial F_1}{\partial z} - \frac{\partial F_3}{\partial x}, \frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y}\right\rangle,$$

where $X = \langle F_1, F_2, F_3\rangle$.

Remark. Given a vector field $X = \langle F_1, F_2\rangle$ on some subset of $\mathbb{R}^2$, we could compute the curl of X by extending X to a (z-invariant) vector field on a subset of $\mathbb{R}^3$:

$$\mathrm{curl}X = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \partial/\partial x & \partial/\partial y & \partial/\partial z \\ F_1 & F_2 & 0 \end{vmatrix} = \left\langle 0, 0, \frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y}\right\rangle.$$

Notice that $\mathrm{curl}X$ is perpendicular to the xy-plane. We might sometimes write $\mathrm{curl}_z X$ for the scalar quantity $\frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y}$, so that $\mathrm{curl}X = (\mathrm{curl}_z X)\mathbf{k}$.

We will try to give a better explanation/motivation for the curl in future classes (such as explaining what this has to do with *curling*). For now, we just use the curl as a necessary condition for conservative vector fields in $\mathbb{R}^3$.

Proposition 4.4. *Let $X : A \to \mathbb{R}^3$ be a C^1-vector field on an open subset A of $\mathbb{R}^3$. If X is conservative, then* $\operatorname{curl} X \equiv \mathbf{0}$ *on A.*

Remark. The notation $\equiv$ is often used for equality between functions, and indicates that the two functions give the same output for every input on the stated domain.

Example 4.5. We can use Proposition 4.4 to show that $X = 4y\mathbf{i} + x\mathbf{j} + 2z\mathbf{k}$ is not conservative. We have

$$F_1 = 4y, \quad F_2 = x, \quad \text{and} \quad F_3 = 2z,$$

so

$$\left\langle \frac{\partial F_3}{\partial y} - \frac{\partial F_2}{\partial z}, \frac{\partial F_1}{\partial z} - \frac{\partial F_3}{\partial x}, \frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y} \right\rangle = \langle 0-0, 0-0, 1-4 \rangle = -3\mathbf{k}.$$

Since this vector field is nonzero, X is not conservative.

Example 4.6. On the other hand, one can check that $\operatorname{curl} X$ does vanish for $X = \langle yz + y, xz + x, xy + 1 \rangle$. We can try to find a scalar potential for X. We want $\phi : \mathbb{R}^3 \to \mathbb{R}$ such that

$$\phi_x = yz + y, \quad \phi_y = xz + x, \quad \text{and} \quad \phi_z = xy + 1.$$

Integrating the first equation, we find that $\phi(x, y, z) = xyz + xy + g(y, z)$, for some function $g(y, z)$. Differentiating this expression with respect to y tells us that

$$xz + x = \phi_y = xz + x + g_y.$$

So $g_y = 0$, meaning that $g(y, z) = h(z)$, for some function h. That is, $\phi(x, y, z) = xyz + xy + h(z)$. Finally, we differentiate this expression with respect to z to find that

$$xy + 1 = \phi_z = xy + h'(z).$$

So $h'(z) = 1$, meaning that $h(z) = z + C$, for some constant C. Altogether, we find that

$$\phi(x, y, z) = xyz + xy + z + C$$

is a scalar potential for X, where C is any real number.

Let's end this section by mentioning what we *won't* do. We've defined what it means for a vector field X on a subset of $\mathbb{R}^n$ to be conservative, for any $n \geq 1$, and we've used Clairaut's theorem to give a necessary condition for conservativity in the cases $n = 2$ and $n = 3$. We can use the same strategy to give a necessary condition in any dimension, but notice that the bookkeeping becomes somewhat horrendous: for a vector field $X = \langle F_1, \ldots, F_n \rangle$, Clairaut's theorem gives

$$\binom{n}{2} = \frac{n(n-1)}{2}$$

differential equations that the component functions $F_1, \ldots, F_n$ must satisfy. We won't try to write down all of these equations, and will instead focus our attention on the cases $n = 2$ and $n = 3$. In the last portion of the quarter, we'll develop some machinery which will allow us to carry out the general analysis much more elegantly.

4.3 Day 11: Curl and divergence

Goals

By the end of today's class, we should be able to do the following.

1. Define the **divergence** of a vector field on $\mathbb{R}^n$.
2. State a necessary condition for a C^1 vector field on $\mathbb{R}^3$ to admit a vector potential.
3. State **Poincaré's lemma**, which gives a sufficient condition for a vector field to admit a scalar potential or vector potential.

4.3.1 Divergence

In our last meeting, we defined a sort of gross-looking vector operation: the curl. So far, this is just a vector field which records a necessary condition for being conservative. But there must be some reason that this operation is called *curl*. Indeed, we will show in due time that $\operatorname{curl} X$ measures the extent to which the vector field X is rotating: the magnitude of $\operatorname{curl} X$ gives a sense of the scale on which X is rotating, and $\operatorname{curl} X$ is perpendicular to the plane in which X is rotating, so that $\operatorname{curl} X$ is supposed to be a perfectly good descriptor for the rotative properties of X.

For now, this geometric description is wholly unjustified, but the point is that $\operatorname{curl} X$ is a sort of vector-valued derivative for vector fields. Indeed, we compute $\operatorname{curl} X$ with derivatives of the component functions of X:

$$\operatorname{curl}\langle F_1, F_2, F_3\rangle = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \partial/\partial x & \partial/\partial y & \partial/\partial z \\ F_1 & F_2 & F_3 \end{vmatrix} = \left\langle \frac{\partial F_3}{\partial y} - \frac{\partial F_2}{\partial z}, \frac{\partial F_1}{\partial z} - \frac{\partial F_3}{\partial x}, \frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y} \right\rangle.$$

Since we have a derivative, we can look for an antiderivative. That is, given a vector field $Y = \langle G_1, G_2, G_3\rangle$, we want to know whether or not we can write $Y = \operatorname{curl} X$ for some other vector field X. This motivates a definition not unlike that of a scalar potential field.

Definition. Let $Y : A \to \mathbb{R}^3$ be a vector field on a subset A of $\mathbb{R}^3$. A **vector potential** (or **vector potential field**) is a vector field $X : A \to \mathbb{R}^3$ satisfying

$$\operatorname{curl} X \equiv Y$$

on A.

Remark. When a vector field admits a scalar potential, we say that the field is *conservative*. We won't use an analogous word for vector fields which admit vector potentials. The three main physics-y adjectives for vector fields that we'll use are *conservative*, *irrotational*, and *incompressible*. Conservative is the only one of these words that is defined by the existence of an antiderivative.

Exercise 4.7. Vector potentials are far from unique. Suppose that X is a vector potential for Y, and that ϕ is any continuously differentiable scalar function. Show that $X + \nabla\phi$ is also a vector potential for Y. So the role played by constants when finding scalar potentials is played by conservative vector fields when finding vector potentials.

Remark. The non-uniqueness of vector potentials is important in ways that I can only pretend to understand. In classical electromagnetism, the magnetic field $\mathbf{B}$ admits a vector potential $\mathbf{A}$ (called the *magnetic vector potential*). Exercise 4.7 tells us that the vector potential $\mathbf{A}$ is not unique, and thus we have an artificial degree of freedom in our mathematical formulation of reality. This makes classical electrodynamics a **gauge theory**, and dealing with this extra degree of freedom requires a **gauge fixing**.

Now suppose that $Y = \langle G_1, G_2, G_3\rangle$ admits a vector potential $X = \langle F_1, F_2, F_3\rangle$. Then

$$\langle G_1, G_2, G_3\rangle = \left\langle \frac{\partial F_3}{\partial y} - \frac{\partial F_2}{\partial z}, \frac{\partial F_1}{\partial z} - \frac{\partial F_3}{\partial x}, \frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y} \right\rangle,$$

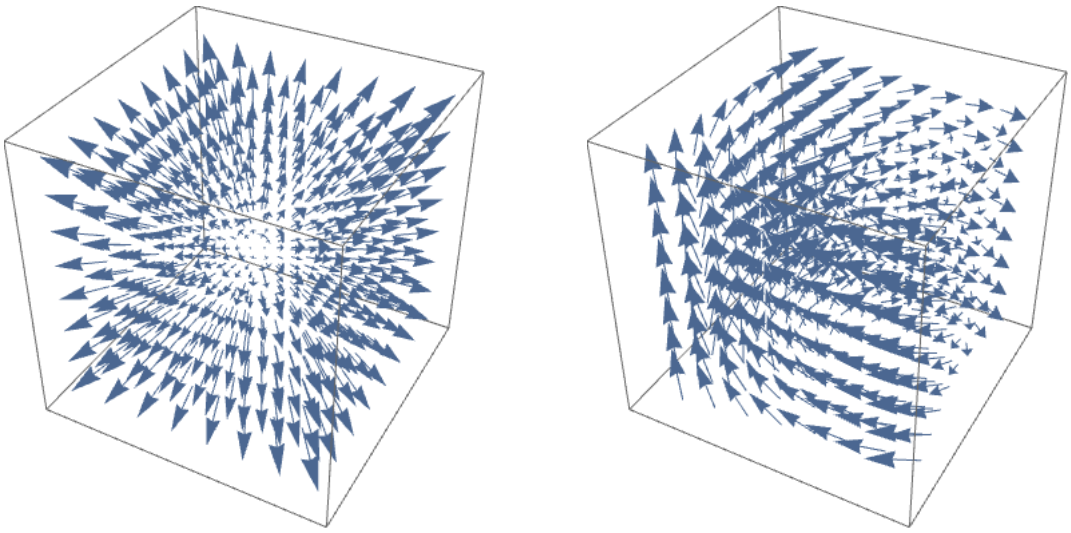

(a) Positive divergence (b) No divergence

Figure 4.2: A vector field which expands, and one which doesn't

and we find that

$$\begin{aligned}\frac{\partial G_1}{\partial x}+\frac{\partial G_2}{\partial y}+\frac{\partial G_3}{\partial z} &= \frac{\partial}{\partial x}\left(\frac{\partial F_3}{\partial y}-\frac{\partial F_2}{\partial z}\right)+\frac{\partial}{\partial y}\left(\frac{\partial F_1}{\partial z}-\frac{\partial F_3}{\partial x}\right)+\frac{\partial}{\partial z}\left(\frac{\partial F_2}{\partial x}-\frac{\partial F_1}{\partial y}\right)\\ &= \frac{\partial^2 F_3}{\partial x\partial y}-\frac{\partial^2 F_2}{\partial x\partial z}+\frac{\partial^2 F_1}{\partial y\partial z}-\frac{\partial^2 F_3}{\partial y\partial x}+\frac{\partial^2 F_2}{\partial z\partial x}-\frac{\partial^2 F_1}{\partial z\partial y}=0,\end{aligned} \tag{4.5}$$

with the last step following from Clairaut's theorem. Another way to write this is via the formal equation

$$\nabla\cdot Y=0,$$

where again $\nabla=\langle\partial/\partial x,\partial/\partial y,\partial/\partial z\rangle$. With this in mind, we define the following scalar-valued derivative of vector fields.

Definition. Let $X:A\to\mathbb{R}^n$ be a C^1-vector field on some open subset A of $\mathbb{R}^n$. The **divergence** of X, denoted $\operatorname{div}X$ or $\nabla\cdot X$, is a scalar-valued function on A defined by

$$\operatorname{div}X:=\nabla\cdot X:=\frac{\partial F_1}{\partial x^1}+\cdots+\frac{\partial F_n}{\partial x^n},$$

where $X=\langle F_1,\ldots,F_n\rangle$.

Remark. Perhaps you could abide the sin of presenting $\operatorname{curl}X$ as just a device for detecting the failure of X to be a gradient. But now we've doubled our transgressions: it can't be the case that both $\operatorname{curl}X$ and $\operatorname{div}X$ are just bookkeeping devices — not with names like that! Indeed, our presentation is rather ahistorical, and it's somewhat painful to keep suppressing the geometric motivation behind these definitions, but some upcoming theorems will unravel all of this for us.

From (4.5) we see that the divergence of a vector field which admits a vector potential is necessarily 0. That is,

$$\operatorname{div}(\operatorname{curl}X)=\nabla\cdot(\nabla\times X)=0.$$

(Notice that the second equality, while purely formal, is kind of nice: the quantity $\nabla\cdot(\nabla\times X)$ looks like a scalar triple product which ought to vanish, since ∇ appears twice.) Based on this, we have a necessary condition for a vector field to admit a vector potential.

Proposition 4.8. *Let $X:A\to\mathbb{R}^3$ be a vector field on an open subset A of $\mathbb{R}^3$. If X admits a vector potential, then $\operatorname{div}X\equiv 0$ on A.*

Example 4.9. From Proposition 4.8 we can easily see that the vector field $X = \langle x, y, z\rangle$, defined on all of $\mathbb{R}^3$ fails to admit a vector potential. Indeed,

$$\operatorname{div} X = \frac{\partial}{\partial x}(x) + \frac{\partial}{\partial y}(y) + \frac{\partial}{\partial z}(z) = 3,$$

and thus we can't have $\operatorname{curl} Y = X$ for any vector field Y. Even though we haven't yet interpreted the divergence geometrically, Figure 4.2a suggests that we've made a meaningful measurement: the vector field X is certainly diverging from the origin.

Exercise 4.10. Based on Example 4.9, produce a vector field whose divergence is everywhere negative. Sketch your plot and make a hypothesis as to what negative divergence tells us.

Example 4.11. On the other hand, if $X = \langle 2y, -2x, 2\rangle$, then

$$\operatorname{div} X = \frac{\partial}{\partial x}(2y) + \frac{\partial}{\partial y}(-2x) + \frac{\partial}{\partial z}(2) = 0,$$

so Proposition 4.8 suggests that X *could* admit a vector potential. Let's see if we can find one. We want

$$Y = \langle G_1, G_2, G_3\rangle$$

so that

$$\left\langle \frac{\partial G_3}{\partial y} - \frac{\partial G_2}{\partial z}, \frac{\partial G_1}{\partial z} - \frac{\partial G_3}{\partial x}, \frac{\partial G_2}{\partial x} - \frac{\partial G_1}{\partial y} \right\rangle = \langle 2y, -2x, 2\rangle. \tag{4.6}$$

To make our computation easier, we exploit the non-uniqueness of vector potentials. Specifically, if Y is a vector potential for X, then so is $Y + \nabla\phi$, for any scalar function ϕ. We can define ϕ to be a z-antiderivative of $-G_3$:

$$\phi(x, y, z) := -\int G_3(x, y, z)\, dz.$$

Then the third component of $\nabla\phi$ is $-G_3$, and thus the third component of $Y + \nabla\phi$ is 0. The upshot is that if X admits a vector potential, then X admits a vector potential whose z-coordinate is identically 0. For this reason we assume that $G_3 \equiv 0$, so that (4.6) reduces to the differential equations

$$\left\langle -\frac{\partial G_2}{\partial z}, \frac{\partial G_1}{\partial z}, \frac{\partial G_2}{\partial x} - \frac{\partial G_1}{\partial y} \right\rangle = \langle 2y, -2x, 2\rangle.$$

Integrating the first two components against z, we find that

$$G_2 = -2yz + g(x, y) \quad \text{and} \quad G_1 = -2xz + h(x, y),$$

for some functions $g(x, y)$ and $h(x, y)$. Then

$$2 = \frac{\partial G_2}{\partial x} - \frac{\partial G_1}{\partial y} = \frac{\partial g}{\partial x} - \frac{\partial h}{\partial y}.$$

This differential equation still leaves us with a good deal of freedom, so we arbitrarily take $h \equiv 0$ and $g(x, y) = 2x + C$, for some constant C. So we see that

$$Y = \langle -2xz, -2yz + 2x, 0\rangle$$

gives a vector potential for X.

Exercise 4.12. Another vector potential for $\langle 2y, -2x, 2\rangle$ is given by $\langle -y, x, x^2 + y^2\rangle$. Prove that, in this case, this potential differs from $\langle -2xz, -2yz + 2x, 0\rangle$ by a conservative vector field. We caution that, in general, this need not be true.

Exercise 4.13. Consider the vector field X, defined on $\mathbb{R}^3 \setminus \{(0, 0, 0)\}$ by

$$X = \frac{1}{(x^2 + y^2 + z^2)^{3/2}} \langle x, y, z\rangle.$$

Show that $\operatorname{div} X = 0$, and then plot X. Your plot should seem to clash with your computation. For a challenge, prove that X does not admit a vector potential. This demonstrates that, while Proposition 4.8 gives a necessary condition for the existence of a vector potential, it does not give a sufficient condition.

4.3.2 Poincaré's lemma

Over the last two class meetings, we've discussed three derivatives which involve vector fields on $\mathbb{R}^3$: the gradient and curl are vector-valued derivatives, while divergence is a derivative which is applied to vector fields. For any open subset A of $\mathbb{R}^3$, these can be assembled into the following diagram:

$$C^\infty(A) \xrightarrow{\text{grad}} \mathfrak{X}(A) \xrightarrow{\text{curl}} \mathfrak{X}(A) \xrightarrow{\text{div}} C^\infty(A) \tag{4.7}$$

Here $C^\infty(A)$ denotes the set of infinitely-differentiable[4] functions $A \to \mathbb{R}$, while $\mathfrak{X}(A)$ denotes the set of smooth vector fields on A. By definition we have

$$\nabla \times \nabla\phi = \mathbf{0} \quad \text{and} \quad \nabla \cdot (\nabla \times X)$$

for any scalar field $\phi : A \to \mathbb{R}$ and any vector field $X : A \to \mathbb{R}^3$, so following two consecutive steps of the diagram in (4.7) gives 0. Another way to write this is as

$$\operatorname{im}\operatorname{grad} \subseteq \ker\operatorname{curl} \quad \text{and} \quad \operatorname{im}\operatorname{curl} \subseteq \ker\operatorname{div}.$$

That is, the **image** of the gradient lies in the **kernel** of curl, and likewise the image of curl lies in the kernel of divergence. The image of gradient is the set of vector fields in $\mathfrak{X}(A)$ which can be written as the gradient of some function, and the kernel of curl is the set of vector fields whose curl is $\mathbf{0}$. A natural question is then: under what circumstances do we have $\operatorname{im}\operatorname{grad} = \ker\operatorname{curl}$ and $\operatorname{im}\operatorname{curl} = \ker\operatorname{div}$? That is, are the necessary conditions given by Propositions 4.2 and 4.4 ever sufficient conditions? Poincaré gives us an affirmative answer.

Theorem 4.14: Poincaré's lemma

Suppose $A \subset \mathbb{R}^3$ is contractible. Then $\operatorname{im}\operatorname{grad} = \ker\operatorname{curl}$ and $\operatorname{im}\operatorname{curl} = \ker\operatorname{div}$.

We won't prove Poincaré's lemma right now, but will return to prove a much more general version later in the quarter. We also won't give a careful definition of *contractible* right now, but will simply say that a subset A of $\mathbb{R}^n$ is **contractible** if it can be continuously shrunk to a point. So, for instance, the punctured Euclidean spaces $\mathbb{R}^2 \setminus \{(0,0)\}$ and $\mathbb{R}^3 \setminus \{(0,0,0)\}$ are not contractible.

Poincaré's lemma explains the phenomena observed in Exercises 4.3 and 4.13. In these exercises we find vector fields which satisfy the necessary conditions of Propositions 4.2 and 4.4, respectively, but still fail to have the desired potential fields. That's because the domains of these vector fields are not contractible, and thus our necessary conditions are not sufficient.

Remark. We point out, however, that a vector field defined on a non-contractible domain *can* have a scalar or vector potential: Poincaré's lemma doesn't rule this out. On non-contractible domains we are (for now) relatively clueless as to what happens.

Example 4.15. Consider the vector field

$$X = \langle -e^{xyz}xy, xy, -xz + e^{xyz}yz \rangle.$$

Without Poincaré's lemma, the only way to be sure that X admits a vector potential is to find one. But since the domain of X is all of $\mathbb{R}^3$ — and is thus contractible — Poincaré's lemma tells us that it's enough to make sure that $\operatorname{div} X = 0$. Indeed,

$$\begin{aligned}\operatorname{div} X &= \frac{\partial}{\partial x}(-e^{xyz}xy) + \frac{\partial}{\partial y}(xy) + \frac{\partial}{\partial z}(-xz + e^{xyz}yz) \\ &= (-e^{xyz}y - e^{xyz}xy^2z) + (x) + (-x + e^{xyz}y + e^{xyz}xy^2z) \\ &= 0.\end{aligned}$$

From Poincaré's lemma, we conclude that X admits a vector potential.

I'm hoping that later this quarter we can investigate the diagram (4.7) more closely for non-contractible subsets of $\mathbb{R}^3$. Poincaré's lemma indicates that this diagram can somehow detect "holes" in our set A, and it turns out that the diagram in (4.7) can detect properties that are even more subtle. Henri Poincaré was a founder of the field which is today known as *topology*, which studies precisely the sorts of properties that (4.7) is able to see.

[4] We don't need ϕ to be infinitely differentiable in order to apply the derivatives in this diagram, but make this assumption just to make our lives easier.

5 Week 5

Last week we discussed two new derivatives for vector fields — curl and div — and discussed some antidifferentiation problems for vector fields. This week we take a look at **line integrals**, which may at first seem wholly unrelated to the derivatives discussed last week. Our week will culminate with **Green's theorem**, which explains how (some) of these line integrals relate to (one) of the vector field derivatives we wrote down last week. This will establish a pattern to be followed for the next few weeks, where we discuss some notion of integration for vector fields, and then relate this integration to a vector derivative.

5.1 Day 12: Line integrals of functions

Goals

By the end of today's class, we should be able to do the following.

1. Define **simple parametrized curves** in $\mathbb{R}^n$.
2. Define the **line integral** of a scalar function along a parametrized curve.
3. Use line integrals to compute **areas** and **masses**.

Consider the disc $D = \{(x, y) : x^2 + y^2 \leq 9\}$ in $\mathbb{R}^2_{x,y}$. Say we want to build a fence around the boundary of this disc, and that the fence's height at the point (x, y) will be given by $h(x, y) = x^2 + 4y^2$. This fence is depicted in Figure 5.1. How might we compute the area of the fence?

We're quite familiar with computing areas via integrals, but since h is a function of two variables, the only integral we have so far is the double integral, which computes volume. Today we will introduce **line integrals**, which, in a possibly confusing twist of terminology, allow us to integrate a function $f : \mathbb{R}^n \to \mathbb{R}$ over a curve $\mathcal{C} \subset \mathbb{R}^n$.

5.1.1 Parametrized curves

In defining line integrals, we will always assume that our curve $\mathcal{C}$ admits a **parametrization**. Moreover, we expect our parametrizations to be injective and differentiable on the interior of their domains. Note that this excludes self-intersecting curves.

Definition. A **simple parametrized curve in $\mathbb{R}^n$** is a piecewise C^1-mapping

$$\mathbf{r}\colon I \to \mathbb{R}^n,$$

where $I \subset \mathbb{R}$ is an interval and $\mathbf{r}$ is injective when restricted to the interior $\mathring{I}$ of I. We say that $\mathbf{r}$ is **regular** if $\mathbf{r}'(t)$ is defined and nonzero, for all $t \in \mathring{I}$.

Remark. We exclude self-intersecting curves from today's discussion simply to make things a little easier; we can reasonably talk about line integrals over such curves if we want.

We will often conflate the **parametrization $\mathbf{r}$** and the **curve $\mathcal{C}$** which is the image of the parametrization, calling them both the curve. This is a little dangerous, however, because a single curve $\mathcal{C} \subset \mathbb{R}^n$ admits infinitely many parametrizations.

Definition. Given a parametrized curve $\mathbf{r}\colon I \to \mathbb{R}^n$, a **reparametrization** of $\mathbf{r}$ is a C^1-mapping

$$\phi : J \to I,$$

where $J \subset \mathbb{R}$ is an interval, such that $\phi'(t) \neq 0$, for all $t \in J$. If $\phi'(t) > 0$ for all $t \in J$, then we say that ϕ is **orientation-preserving**; otherwise we say that ϕ is **orientation-reversing**.

Remark. In another abuse of terminology, we will often refer to the parametrized curve $\mathbf{r} \circ \phi : J \to \mathbb{R}^n$ as a reparametrization of $\mathbf{r}$.

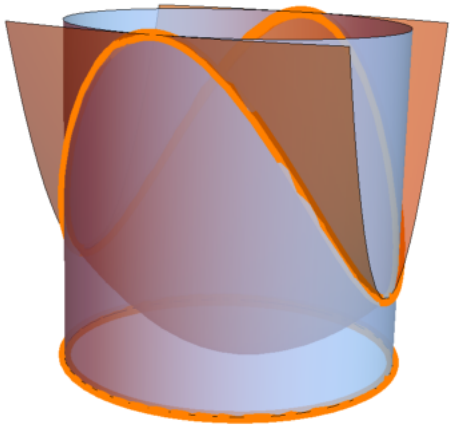

Figure 5.1: We can use line integration to compute the area on the cylinder $x^2 + y^2 = 9$ satisfying $0 \le z \le x^2 + 4y^2$.

Example 5.1. You are likely familiar with computing **arclength reparametrizations** of regular curves. To construct such reparametrizations, we fix a value $t_0 \in I$ and consider the arclength function $g : I \to \mathbb{R}$, defined by

$$g(t) := \int_{t_0}^{t} \|\mathbf{r}'(u)\| \, du.$$

Because $\mathbf{r}$ is regular, $g'(t) = \|\mathbf{r}'(t)\|$ never vanishes, and g is a bijection onto its image $J := g(I)$. We can therefore define $\phi : J \to I$ to be the inverse of g, and we find that

$$\mathbf{r} \circ \phi : J \to \mathbb{R}^n$$

is a unit-speed reparametrization of the curve $\mathcal{C}$ parametrized by $\mathbf{r}$.

5.1.2 Line integrals

We can now give a sketchy description of how one would define line integrals. We won't be particularly careful with this definition, which we rarely use in practice. Throughout, we take

$$\mathbf{r} : I \to \mathbb{R}^2$$

to be a regular, simple parametrized curve, and we take $f : \mathbb{R}^2 \to \mathbb{R}$ to be a continuous function. The story is completely analogous in higher dimensions.

First, let's describe geometrically the area that our integral should compute: we have a surface $\{(\mathbf{r}(t), z) : t \in I, z \in \mathbb{R}\}$, which is the cylinder in $\mathbb{R}^3$ whose cross-section at $z = 0$ is parametrized by $\mathbf{r}$. We're hoping to compute the area of $\mathcal{S}$, which is that portion of our surface bounded by the plane $z = 0$ and the graph $z = f(x, y)$. Again, see Figure 5.1.

As with any integral, we proceed by taking increasingly accurate approximations. We may partition the interval I into subintervals $J_1, \ldots, J_n$ and then ask for the relevant area above each subinterval J_k. One wrinkle we have is that, in addition to having a non-constant height over J_k (as was the case in regular, single-variable integration), the base of our two-dimensional region is itself curved. That is, even if we approximate the value of $f(\mathbf{r}(t))$ over the interval J_k, we can't very well compute the area of a surface with base parametrized by $\mathbf{r} : J_k \to \mathbb{R}^2$. So we must also approximate the base. Namely, if we have $J_k = [t_{k-1}, t_k]$, then we consider the line segment parametrized by

$$\ell_k(\lambda) := (1 - \lambda)\mathbf{r}(t_{k-1}) + \lambda \mathbf{r}(t_k), \quad 0 \le \lambda \le 1.$$

This line segment in $\mathbb{R}^2$ approximates the image of $\mathbf{r} : J_k \to \mathbb{R}^2$, and over this line segment we may approximate the height of our surface $\mathcal{S}$ by $f(t_{k-1})$, which is the height at the left endpoint t_{k-1}. See Figure 5.2.

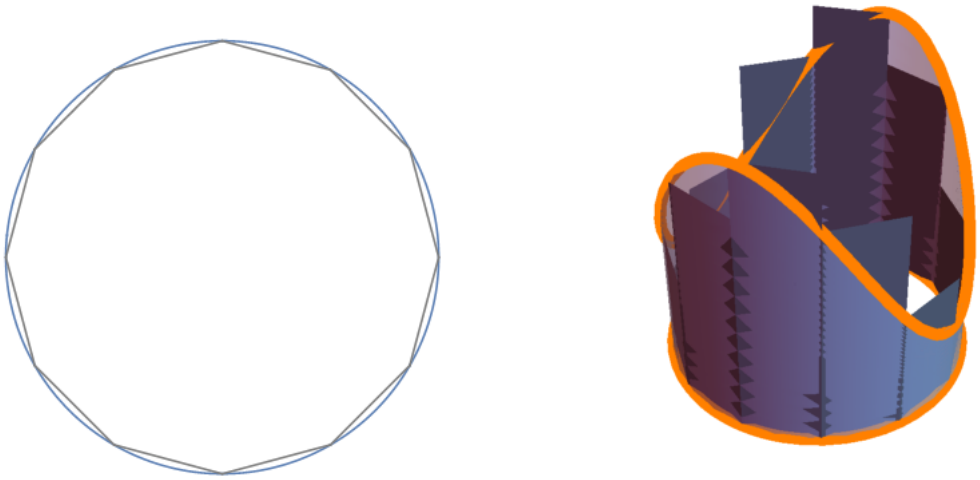

(a) Approximating $\mathcal{C}$ (b) Approximating $\mathcal{S}$

Figure 5.2: When defining line integrals, we approximate the surface over $\mathcal{C}$ by approximating f with a piecewise constant function and approximating $\mathcal{C}$ with a polygon.

Hopefully we can see from this vague description that the area in which we're interested should be approximated by

$$\text{area}(\mathcal{S}) \approx \sum_{k=1}^{n} f(\mathbf{r}(t_{k-1}))\|\mathbf{r}(t_k) - \mathbf{r}(t_{k-1})\|, \tag{5.1}$$

since each of the terms in our sum is the height of one of our rectangles, multiplied by the length of its base. We will very loosely define the **line integral** of f over the curve $\mathcal{C}$ which is parametrized by $\mathbf{r}$ to be the limit of such approximations as our partitions become arbitrarily fine:

$$\int_{\mathcal{C}} f \, ds := \lim_{n\to\infty} \sum_{k=1}^{n} f(\mathbf{r}(t_{k-1}))\|\mathbf{r}(t_k) - \mathbf{r}(t_{k-1})\|.$$

Of course this is very far from being rigorous, since (a) the notation t_k isn't really well-defined as n varies; and (b) we haven't addressed the question of which functions are integrable over $\mathcal{C}$. For this reason, we don't give the line integral the honor of a boxed definition.

Remark. Even though we still haven't given a rigorous treatment of differential forms, take note of the notation we use for line integrals. We integrate against the differential form ds, which measures arclength. This doesn't have any actual meaning at this point in the course, but one takeaway is that we don't want the line integral of f over $\mathcal{C}$ to depend on how we parametrize $\mathcal{C}$. Since arclength is independent of parametrization, this is a reasonable quantity against which to integrate.

At this point, line integrals continue following the usual story for integrals: we write down a definition in terms of partitions, and then we find a way to not use that definition directly. Our first observation is that if we have a very fine partition, then t_k is not that much larger than t_{k-1}, and thus we can make a local linear approximation:

$$\mathbf{r}(t_k) \approx \mathbf{r}(t_{k-1}) + (t_k - t_{k-1})\mathbf{r}'(t_{k-1}).$$

So we see that

$$\|\mathbf{r}(t_k) - \mathbf{r}(t_{k-1})\| \approx \|(t_k - t_{k-1})\mathbf{r}'(t_{k-1})\| = (t_k - t_{k-1})\|\mathbf{r}'(t_{k-1})\|,$$

and thus

$$\text{area}(\mathcal{S}) \approx \sum_{k=1}^{n} f(\mathbf{r}(t_{k-1}))\|\mathbf{r}'(t_{k-1})\|(t_k - t_{k-1}).$$

But this should look familiar. Specifically, this is the nth left-hand Riemann sum for the function $g : I \to \mathbb{R}$ defined by

$$g(t) := f(\mathbf{r}(t))\|\mathbf{r}'(t)\|,$$

where we compute the Riemann sum using the partition $\mathcal{P} = \{t_0, t_1, \ldots, t_n\}$. So we've given a not-very-rigorous justification for the following result.

Theorem 5.2: Computing scalar line integrals

Let $\mathbf{r}\colon I \to \mathbb{R}^n$ be a regular, simple parametrization of a curve $\mathcal{C} \subset \mathbb{R}^n$. If $f : \mathbb{R}^n \to \mathbb{R}$ is continuous, then

$$\int_{\mathcal{C}} f \, ds = \int_I f(\mathbf{r}(t)) \|\mathbf{r}'(t)\| \, dt.$$

Proving this result carefully is kind of tedious (not unlike proving the change of variables theorem), so we won't do it. But if you already believe the single-variable change of variables theorem, then we can justify this theorem in a different way. Consider the following exercise.

Exercise 5.3. Show that if $\mathbf{r}\colon I \to \mathbb{R}^n$ is a regular, simple parametrized curve and $\phi : J \to I$ is a reparametrization, then

$$\int_J f((\mathbf{r} \circ \phi)(t)) \left\|(\mathbf{r} \circ \phi)'(t)\right\| \, dt = \int_I f(\mathbf{r}(t)) \|\mathbf{r}'(t)\| \, dt.$$

So according to Exercise 5.3, we could define the line integral of f along $\mathcal{C}$ by taking *any* regular, simple parametrization of $\mathcal{C}$ and computing the integral on the right hand side of Theorem 5.2. Exercise 5.3 guarantees that this definition makes sense.

Remark. We will sometimes use the notation $\oint_{\mathcal{C}} f \, ds$ for line integrals, in the event that $\mathcal{C}$ is a closed curve. For us, this doesn't have any special meaning, except to point out that $\mathcal{C}$ is a closed curve. I will probably not remember to use the notation $\oint_{\mathcal{C}} f \, ds$, which further specifies that $\mathcal{C}$ is oriented counterclockwise.

5.1.3 A slightly different take

A different (but also not rigorous) way of thinking about line integrals focuses on the similarities between Theorem 5.2 and the change of variables theorem. Diagrammatically we have

$$\begin{array}{ccc} \mathbb{R}^n & \xrightarrow{\Phi} & \mathbb{R}^n \\ & & \downarrow f \\ & & \mathbb{R} \end{array} \qquad\qquad \begin{array}{ccc} \mathbb{R} & \xrightarrow{\mathbf{r}} & \mathbb{R}^n \\ & & \downarrow f \\ & & \mathbb{R} \end{array}. \tag{5.2}$$

On the left is the setup for the change of variables theorem. We have a function $f : \mathbb{R}^n \to \mathbb{R}$ which we want to integrate over some (n-dimensional) domain in $\mathbb{R}^n$, with coordinates $x^1, \ldots, x^n$, and we do this by pulling our domain back to a (hopefully nicer-looking) domain in the copy of $\mathbb{R}^n$ which has coordinates $u^1, \ldots, u^n$. As a result, we pick up a "fudge factor" of $|\det \Phi_*|$ — this measures the extent to which Φ distorts volumes.

On the right side of (5.2) we depict the setup for line integrals. Once again we have a function $f : \mathbb{R}^n \to \mathbb{R}$ which we want to integrate, but this time we want to integrate over some curve $\mathcal{C} \subset \mathbb{R}^n$. In order to do this, we pull our integral back under some parametrization $\mathbf{r}\colon I \subset \mathbb{R} \to \mathcal{C}$. This has the effect of straightening out our curve, but comes at the expense of picking up another fudge factor. This time, lengths are distorted by a factor of $\|\mathbf{r}'(t)\|$, so this is the term we pick up in our integral. Notice that this is an agreement with our existing understanding of arclength, since it tells us that

$$\int_{\mathcal{C}} 1 \, ds = \int_I \|\mathbf{r}'(t)\| \, dt,$$

and surely integrating the constant function 1 should give arclength.

The approach diagrammed in (5.2) will be used next week to define surface integration, and in a more general setting we can use it to define integration on smooth manifolds.

5.1.4 Examples

As a first example, we can solve the problem which motivated this section.

Example 5.4. We want to find the area of the surface $\mathcal{S}$ determined by the equation $x^2+y^2=9$ and the inequality $0\le z\le x^2+4y^2$ in $\mathbb{R}^3_{x,y,z}$. This area will be given by

$$\int_{\mathcal{C}} h\,ds,$$

where $\mathcal{C}$ is the circle $\{(x,y): x^2+y^2=9\}$ in $\mathbb{R}^2$, and $h(x,y)=x^2+4y^2$. We want to use Theorem 5.2 to compute this line integral, so we need a regular parametrization of $\mathcal{C}$. Certainly $\mathbf{r}\colon[0,2\pi]\to\mathbb{R}^2$, defined by

$$\mathbf{r}(t):=(3\cos t,3\sin t),$$

will do the trick. Notice that $\mathbf{r}$ is C^1, since $\mathbf{r}'(t)$ is defined for all $t\in(0,2\pi)$, and $\mathbf{r}$ is injective when restricted to $(0,2\pi)$. According to Theorem 5.2, we have

$$\begin{aligned}\int_{\mathcal{C}} h\,ds &= \int_0^{2\pi} h(\mathbf{r}(t))\|\mathbf{r}'(t)\|\,dt = \int_0^{2\pi} h(3\cos t,3\sin t)\|\langle -3\sin t,3\cos t\|\,dt\\ &= \int_0^{2\pi}(9\cos^2 t+36\sin^2 t)(3)\,dt = 27\int_0^{2\pi}(1+3\sin^2 t)\,dt\\ &= 27\int_0^{2\pi}\left(1+3\frac{1-\cos(2t)}{2}\right)dt = 27\left[t+\frac{3}{2}(t-\frac{1}{2}\sin(2t))\right]_0^{2\pi}\\ &= 27[2\pi+3\pi]=\boxed{135\pi}.\end{aligned}$$

In addition to computing areas, we can also apply line integrals in many of the same ways that we apply our usual, single-variable integrals.

Example 5.5. Consider a metal wire whose shape is parametrized by $\mathbf{r}(t)=(\cos t,\sin t,t^2)$, $t\in[0,2\pi]$. Given that the mass density function for this wire is $\rho(x,y,z)=\sqrt{z}$, let's find the mass of the wire. As usual, we can calculate mass by integrating density:

$$M=\int_{\mathcal{C}}\rho\,ds,$$

where $\mathcal{C}$ is the curve traced out by the wire. So, according to Theorem 5.2, we have

$$\begin{aligned}M &= \int_0^{2\pi}\rho(\mathbf{r}(t))\|\mathbf{r}'(t)\|\,dt = \int_0^{2\pi}\sqrt{t^2}\|\langle-\sin t,\cos t,2t\rangle\|\,dt\\ &= \int_0^{2\pi}|t|\sqrt{1+4t^2}\,dt = \frac{1}{8}\int_1^{1+16\pi^2}\sqrt{u}\,du\\ &= \frac{1}{8}\left[\frac{2}{3}u^{3/2}\right]_1^{1+16\pi^2}=\boxed{\frac{1}{12}((1+16\pi^2)^{3/2}-1)}.\end{aligned}$$

Of course this closed form probably isn't that useful to a physicist or engineer, but it works for a mathematician. The mathematician is also a bit lazy about units: ρ is measured in g/cm, so M is given in g.

5.2 Day 13: Line integrals of vector fields and flux

Goals

By the end of today's class, we should be able to do the following.

1. Define and interpret the **line integral** of a vector field along a parametrized curve.
2. Compute the **work** done by a vector field along a parametrized curve, as well as the **flux** of a vector field through a parametrized curve in $\mathbb{R}^2$.

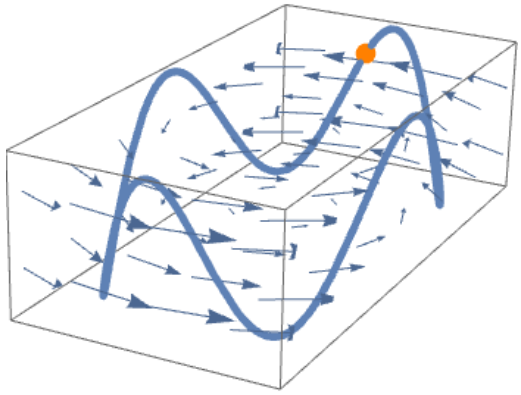

Figure 5.3: We model a wire submerged in flowing water with a curve $\mathcal{C}$ and a vector field $\mathbf{F}$. The orange bead is moved along the wire by the water.

5.2.1 Work

Consider a metal wire which is completely submerged in moving water, and suppose there is a bead on the wire, free to move along the wire. We can model this situation by treating the wire as a curve $\mathcal{C} \subset \mathbb{R}^3$ and representing the force of the flowing water as a vector field $\mathbf{F}$ on $\mathbb{R}^3$. See Figure 5.3.

We'd like to know how much **work** the vector field $\mathbf{F}$ does along $\mathcal{C}$. That is, if we want to move the bead along the wire, how much help do we get from $\mathbf{F}$?

Our first step towards answering this question is to determine the component of $\mathbf{F}$ along $\mathcal{C}$. That is, how much is $\mathbf{F}$ pushing along $\mathcal{C}$ at a point P on $\mathcal{C}$? Notice that, in order for this question to make sense, *we need $\mathcal{C}$ to be an oriented curve*. To this end, let's assume that $\mathbf{r}\colon I \to \mathbb{R}^2$ is a regular, simple parametrization of $\mathcal{C}$. This choice of parametrization determines an orientation on $\mathcal{C}$, in that we will take the positive tangent direction at $\mathbf{r}(t)$ to be the direction in which $\mathbf{r}'(t)$ is pointing. In particular, choosing an orientation for $\mathcal{C}$ determines a unit tangent vector $\mathbf{T}$ to $\mathcal{C}$ at each point P of $\mathcal{C}$. The component of $\mathbf{F}$ along $\mathbf{T}$ is then simply $\mathbf{F} \cdot \mathbf{T}$. See Figure 5.4a. We calculate the total work done by $\mathbf{F}$ along $\mathcal{C}$ by integrating this quantity over all of $\mathcal{C}$.

Definition. Let $\mathcal{C} \subset \mathbb{R}^n$ be an oriented curve which admits a regular, simple parametrization, and let $\mathbf{F}\colon A \to \mathbb{R}^n$ be a vector field defined on some subset A of $\mathbb{R}^n$ containing $\mathcal{C}$. The **line integral of $\mathbf{F}$ along $\mathcal{C}$** is defined by

$$\int_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} := \int_{\mathcal{C}} (\mathbf{F} \cdot \mathbf{T})\, ds, \tag{5.3}$$

where $\mathbf{T}$ is the oriented unit tangent vector to $\mathcal{C}$, provided the integral on the right exists. The quantity $\int_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r}$ is the **work** done by $\mathbf{F}$ along $\mathcal{C}$.

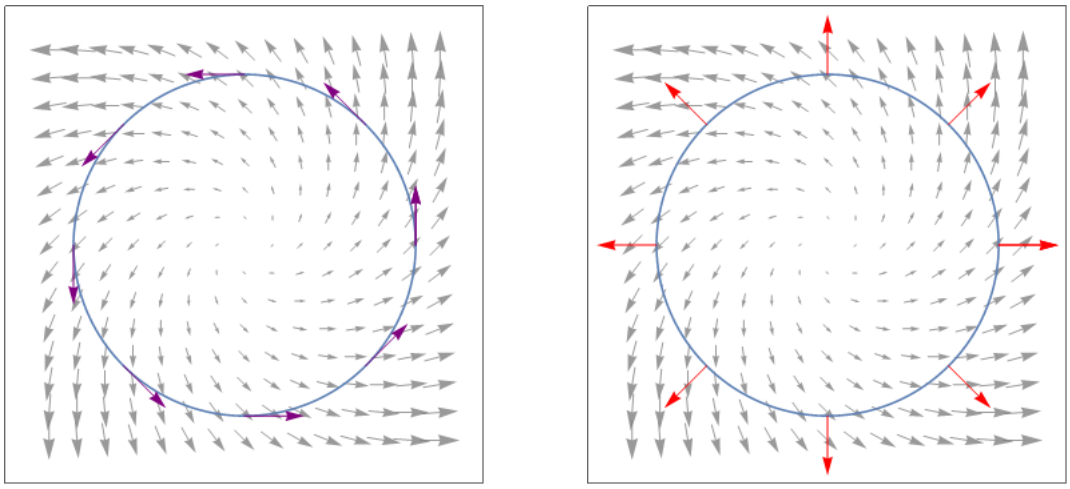

(a) Work (b) Flux

Figure 5.4: We can measure the *work* of $\mathbf{F}$ along $\mathcal{C}$ by integrating the dot product $\mathbf{F}\cdot\mathbf{T}$, and measure the *flux* of $\mathbf{F}$ through $\mathcal{C}$ by integrating $\mathbf{F}\cdot\mathbf{N}$.

Remark.

1. Notice that (5.3) defines the line integral of a vector field to be the line integral of a certain scalar function determined by the vector field.
2. Also notice that, while the discussion preceding our definition focused on curves in $\mathbb{R}^2$ and $\mathbb{R}^3$, we define the line integral of a vector field in any dimension.

Now suppose, as above, that we have a regular, simple parametrization $\mathbf{r}\colon I\to\mathbb{R}^n$ of $\mathcal{C}\subset\mathbb{R}^n$. Moreover, suppose that $\mathbf{r}$ is an **oriented** parametrization of $\mathcal{C}$, meaning that $\mathbf{r}'(t)$ points in the same direction as $\mathbf{T}$, for all $t\in I$. Then

$$\mathbf{T}=\frac{\mathbf{r}'(t)}{\|\mathbf{r}'(t)\|}.$$

So we can apply Theorem 5.2 to Equation 5.3 to conclude that

$$\boxed{\int_{\mathcal{C}}\mathbf{F}\cdot d\mathbf{r}=\int_I\left(\mathbf{F}(\mathbf{r}(t))\cdot\frac{\mathbf{r}'(t)}{\|\mathbf{r}'(t)\|}\right)\|\mathbf{r}'(t)\|\,dt=\int_I\mathbf{F}(\mathbf{r}(t))\cdot\mathbf{r}'(t)\,dt}.$$

So we'll find that vector line integrals are often less cumbersome to compute than scalar line integrals, since the term $\|\mathbf{r}'(t)\|$ need not make an appearance.

5.2.2 Examples

Example 5.6. Consider the vector field $\mathbf{F}\colon\mathbb{R}^2\to\mathbb{R}^2$ defined by $\mathbf{F}(x,y)=\langle\cos x,\sin x\rangle$ and the curve $\mathcal{C}$ parametrized by $\mathbf{r}(t)=(t,t^2)$, $-\pi\le t\le\pi/2$. We have

$$\begin{aligned}\int_{\mathcal{C}}\mathbf{F}\cdot\,d\mathbf{r}&=\int_{-\pi}^{\pi/2}\mathbf{F}(\mathbf{r}(t))\cdot\mathbf{r}'(t)\,dt=\int_{-\pi}^{\pi/2}\langle\cos t,\sin t\rangle\cdot\langle1,2t\rangle\,dt\\&=\int_{-\pi}^{\pi/2}(\cos t+2t\sin t)\,dt=[-2t\cos t+3\sin t]_{-\pi}^{\pi/2}\\&=[(-\pi\cos(\pi/2)+3\sin(\pi/2))-(2\pi\cos(-\pi)+3\sin(-\pi))]=\boxed{3+2\pi}.\end{aligned}$$

Remark. There's another notation for vector line integrals which is sufficiently popular (especially among physicists and engineers) to merit discussion. If $\mathbf{F}=\langle F_1,F_2\rangle$, we will sometimes write the line integral of $\mathbf{F}$

over $\mathcal{C} \subset \mathbb{R}^2$ as

$$\int_{\mathcal{C}} F_1(x,y)\,dx + F_2(x,y)\,dy, \tag{5.4}$$

and similarly for vector line integrals in dimension three. One reason for the popularity of this notation is that we can somehow sidestep any thoughts of dot products. Namely, if $\mathbf{r}(t) = (x(t), y(t))$, then $\mathbf{r}'(t) = \langle x'(t), y'(t)\rangle$, so

$$\mathbf{F}(\mathbf{r}(t)) \cdot \mathbf{r}'(t) = \langle F_1(x(t),y(t)), F_2(x(t),y(t))\rangle \cdot \langle x'(t), y'(t)\rangle = F_1(x(t),y(t))x'(t) + F_2(x(t),y(t)).$$

From this we see that

$$\int_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = \int_I (F_1(x(t),y(t))x'(t) + F_2(x(t),y(t))y'(t))\,dt,$$

which is precisely what we would get from (5.4) by naïvely replacing x and y with $x(t)$ and $y(t)$.

Example 5.7. Let's compute the value of

$$I = \int_{ABC} \frac{x\,dy - y\,dx}{x^2+y^2},$$

where $A = (-1,0)$, $B = (0,1)$, $C = (1,0)$, and ABC is the piecewise linear path running from A to B to C. Our first observation is that

$$I = \int_{AB} \frac{x\,dy - y\,dx}{x^2+y^2} + \int_{BC} \frac{x\,dy - y\,dx}{x^2+y^2},$$

so we consider parametrizations

$$\mathbf{r}_{AB}(t) = (t-1, t), \quad 0 \le t \le 1 \quad \text{and} \quad \mathbf{r}_{BC}(t) = (t, 1-t), \quad 0 \le t \le 1$$

for AB and BC, respectively. Now $I = \int_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r}$, where

$$\mathbf{F}(x,y) = \left\langle -\frac{y}{x^2+y^2}, \frac{x}{x^2+y^2} \right\rangle,$$

so

$$\mathbf{F}(\mathbf{r}_{AB}(t)) = \left\langle -\frac{t}{(t-1)^2+t^2}, \frac{t-1}{(t-1)^2+t^2} \right\rangle \quad \text{and} \quad \mathbf{F}(\mathbf{r}_{BC}(t)) = \left\langle -\frac{1-t}{t^2+(1-t)^2}, \frac{t}{t^2+(1-t)^2} \right\rangle.$$

Finally, $\mathbf{r}'_{AB}(t) = \langle 1, 1\rangle$ and $\mathbf{r}'_{BC}(t) = \langle 1, -1\rangle$, so

$$\begin{aligned}
\int_{AB} \frac{x\,dy - y\,dx}{x^2+y^2} &= \int_0^1 \mathbf{F}(\mathbf{r}_{AB}(t)) \cdot \mathbf{r}'_{AB}(t)\,dt = \int_0^1 \frac{(t-1)-t}{(t-1)^2+t^2}\,dt \\
&= -\int_0^1 \frac{1}{2t^2-2t+1}\,dt = -\int_0^1 \frac{2}{(2t-1)^2+1}\,dt \\
&= [-\arctan(2t-1)]_0^1 = \arctan(-1) - \arctan(1) = -\frac{\pi}{2}.
\end{aligned}$$

Similarly,

$$\begin{aligned}
\int_{BC} \frac{x\,dy - y\,dx}{x^2+y^2} &= \int_0^1 \mathbf{F}(\mathbf{r}_{BC}(t)) \cdot \mathbf{r}'_{BC}(t)\,dt = \int_0^1 -\frac{(1-t)+t}{t^2+(1-t)^2}\,dt \\
&= -\int_0^1 \frac{1}{2t^2-2t+1}\,dt = -\frac{\pi}{2},
\end{aligned}$$

with the last step using the earlier computations. Altogether, we find that $\boxed{I = -\pi}$.

It's worth pointing out that approaching these integrals like a physicist makes them a bit easier. Namely, a physicist would note that AB has the equation $y = x+1$, and thus $dy = dx$. So she'd simply write

$$\int_{AB} \frac{x\,dy - y\,dx}{x^2+y^2} = \int_{-1}^0 \frac{x\,dx - (x+1)\,dx}{x^2+(x+1)^2} = \cdots = -\frac{\pi}{2},$$

with the bounds of integration determined by the range of values attained by x. Similarly, BC has equation $y = 1-x$, so $dy = -dx$, and thus

$$\int_{BC} \frac{x\,dy - y\,dx}{x^2+y^2} = \int_0^1 \frac{-x\,dx - (1-x)\,dx}{x^2+(1-x)^2} = \cdots = -\frac{\pi}{2}.$$

This is a decent way to check our work, but I'm afraid that in this class we'll be too pedantic to find this an acceptable argument. Our main objection is that phrases like $dy = -dx$ don't have any meaning to us yet. Once we've defined integration of differential forms, these sorts of substitutions will be acceptable.

5.2.3 Flux

From Figure 5.4b we see that if $\mathcal{C}$ is a curve in $\mathbf{R}^2$, then $\mathbf{F}\cdot\mathbf{T}$ is not the only reasonable dot product to consider along $\mathcal{C}$. While integrating $\mathbf{F}\cdot\mathbf{T}$ tells us the total work done by $\mathbf{F}$ along $\mathcal{C}$, we can integrate $\mathbf{F}\cdot\mathbf{N}$ over $\mathcal{C}$ to compute the total **flux** of $\mathbf{F}$ through $\mathcal{C}$, where $\mathbf{N}$ is the **oriented unit normal vector** to $\mathcal{C}$. The total flux of $\mathbf{F}$ through $\mathcal{C}$ intends to measure the extent to which $\mathbf{F}$ is pushing out of $\mathcal{C}$.

Remark. We're being pretty vague here about precisely what flux measures. This is somewhat intentional. There are many forms of flux one might want to compute (volumetric flux, heat flux, energy flux, momentum flux,...), and we're trying to develop a mathematical model which is amenable to any of these. In general, flux through a curve has dimensions [quantity]/([time]$\cdot$[length]), while flux through a surface has dimensions [quantity]/([time]$\cdot$[area]), where *quantity* depends on the flux of interest to us.

Remark. Note that we can only compute the flux of $\mathbf{F}$ through $\mathcal{C}$ if $\mathcal{C}$ is a curve in $\mathbb{R}^2$. In general, we compute the flux of a vector field through an object only when that object has one fewer dimension than the ambient space.

Just as with the line integral of a vector field, it would be nice if we could write down a formula for computing the total flux of $\mathbf{F}$ through $\mathcal{C}$ in terms of a parametrization of $\mathbf{r}(t)$ of $\mathcal{C}$. Indeed we can. Let's assume that $\mathbf{r}(t) = (x(t), y(t))$, $a \le t \le b$, is an oriented parametrization of $\mathcal{C}$. Then

$$\mathbf{r}'(t) = \langle x'(t), y'(t)\rangle$$

gives a properly tangent vector to $\mathcal{C}$. We can obtain a normal vector to $\mathcal{C}$ by rotating $\mathbf{r}'(t)$ clockwise through an angle of $\pi/2$ radians[1]. This produces a normal vector of $\langle y'(t), -x'(t)\rangle$, and we can normalize to obtain a unit normal vector:

$$\mathbf{N}(\mathbf{r}(t)) = \frac{\langle y'(t), -x'(t)\rangle}{\|\mathbf{r}'(t)\|}.$$

We can now compute the flux:

$$\begin{aligned}\int_{\mathcal{C}} (\mathbf{F}\cdot\mathbf{N})\,ds &= \int_a^b \left(\mathbf{F}(\mathbf{r}(t)) \cdot \frac{\langle y'(t), -x'(t)\rangle}{\|\mathbf{r}'(t)\|}\right)\|\mathbf{r}'(t)\|\,dt \\ &= \int_a^b \langle F_1(x(t),y(t)), F_2(x(t),y(t))\rangle \cdot \langle y'(t), -x'(t)\rangle\,dt.\end{aligned}$$

Maybe this isn't too attractive, but it gives us a means of computing the flux of $\mathbf{F}$ through $\mathcal{C}$ with a given parametrization, and also allows us to relate flux to a line integral of a vector field.

Exercise 5.8. Suppose $\mathbf{F} = \langle F_1, F_2\rangle$, and that the oriented unit normal vector $\mathbf{N}$ of $\mathcal{C}$ is obtained from the oriented unit tangent vector $\mathbf{T}$ via a clockwise rotation. Prove that the flux of $\mathbf{F}$ through $\mathcal{C}$ is given by $\int_{\mathcal{C}} \widehat{\mathbf{F}}\cdot d\mathbf{r}$, where $\widehat{\mathbf{F}} = \langle -F_2, F_1\rangle$.

We make the important note that *flux is not defined as a vector line integral*. Instead, we compute flux via a particular scalar line integral. By their definition, vector line integrals compute work.

Example 5.9. Figure 5.5 depicts the curve $\mathcal{C}$ determined by $x^2 + 4y^2 = 16$ in $\mathbb{R}^2$, as well as the vector fields

$$\mathbf{F}_1 = \langle 0, 1\rangle, \quad \mathbf{F}_2 = \langle x, -y\rangle, \quad \text{and} \quad \mathbf{F}_3 = \langle x, y\rangle.$$

[1] Of course we could also get a normal vector via counterclockwise rotation. Which normal vector we prefer will depend on the situation, but if our curve has a counterclockwise orientation and an outward-pointing normal vector, then the convention we use here will do the trick.

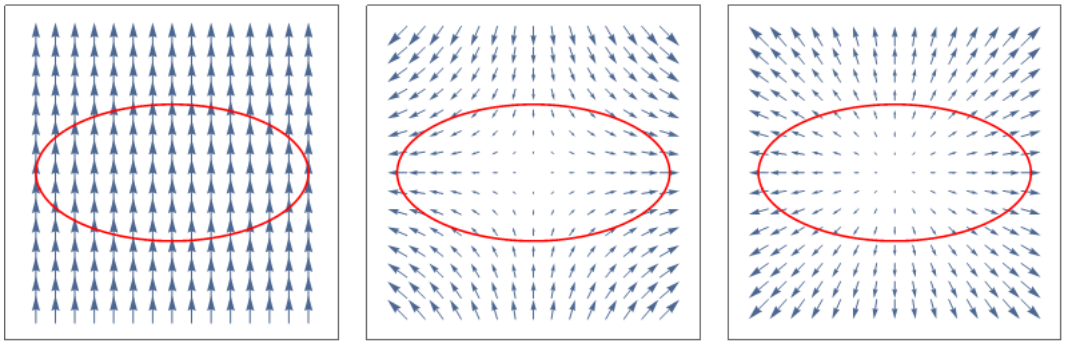

Figure 5.5: Various vector fields flowing through a region bounded by $\mathcal{C} = \{(x, y) : x^2 + 4y^2 = 16\}$.

In each case, the total flux of $\mathbf{F}_i$ through $\mathcal{C}$ is given by

$$\int_{\mathcal{C}} (\mathbf{F}_i \cdot \mathbf{N})\, ds.$$

We can see that the flux of $\mathbf{F}_1$ through $\mathcal{C}$ ought to be 0, while that of $\mathbf{F}_3$ should be positive. The flux of $\mathbf{F}_2$ through $\mathcal{C}$ is less clear.

Let's compute the flux of $\mathbf{F}_k$ through $\mathcal{C}$ for $k = 1, 2, 3$, where we take $\mathcal{C}$ to have an outward-pointing normal vector. First, we parametrize $\mathcal{C}$ via

$$\mathbf{r}(t) = (4\cos t, 2\sin t), \quad 0 \le t \le 2\pi.$$

So $x(t) = 4\cos t$ and $y(t) = 2\sin t$, and thus $x'(t) = -4\sin t$, while $y'(t) = 2\cos t$. Consider the composition of $\mathbf{F}_k$ with $\mathbf{r}(t)$, for $k = 1, 2, 3$:

$$\mathbf{F}_1(\mathbf{r}(t)) = \langle 0, 1\rangle, \quad \mathbf{F}_2(\mathbf{r}(t)) = \langle 4\cos t, -2\sin t\rangle, \quad \text{and} \quad \mathbf{F}_3(\mathbf{r}(t)) = \langle 4\cos t, 2\sin t\rangle.$$

Then the flux of $\mathbf{F}_1$ through $\mathcal{C}$ is

$$\int_0^{2\pi} \langle 0, 1\rangle \cdot \langle 2\cos t, 4\sin t\rangle \, dt = \int_0^{2\pi} 4\sin t \, dt = 0;$$

the flux of $\mathbf{F}_2$ through $\mathcal{C}$ is

$$\int_0^{2\pi} \langle 4\cos t, -2\sin t\rangle \cdot \langle 2\cos t, 4\sin t\rangle \, dt = \int_0^{2\pi} (8\cos^2 t - 8\sin^2 t) \, dt = \int_0^{2\pi} 8\cos(2t) \, dt = 0;$$

and the flux of $\mathbf{F}_3$ through $\mathcal{C}$ is

$$\int_0^{2\pi} \langle 4\cos t, 2\sin t\rangle \cdot \langle 2\cos t, 4\sin t\rangle \, dt = \int_0^{2\pi} 8 \, dt = 16\pi.$$

The first and third computations match our hypotheses; the middle computation tells us that the flow exiting our region through the left and right is balanced out by inward flow along the top and bottom.

5.3 Day 14: Green's theorems for work and flux

Goals

By the end of today's class, we should be able to do the following.

1. Use **Green's theorem** to relate the work done by a vector field along a plane curve to the curl of the vector field.
2. Use the flux version of Green's theorem to relate the flux of a vector field through a plane curve to the divergence of the vector field.

Today we will relate the line integrals we've defined this week to the derivatives we considered last week. The primary result is **Green's theorem**, which is a two-dimensional version of the Fundamental Theorem of Integral Calculus. We will also discuss a flux version of Green's theorem, and use these two results to give a geometric justification for calling last week's derivatives *divergence* and *curl*.

5.3.1 Green's theorem

As mentioned above, Green's theorem will relate line integrals to a vector derivative. Somewhat more specifically, we will integrate the function $\operatorname{curl}_z(\mathbf{F})$ over a region, and compare this integral to a line integral along the boundary of the region. You'll recall that vector line integrals require the underlying curve to be oriented, so we establish a convention for our boundary orientations.

Definition. Suppose $\mathcal{D}$ is a closed, bounded region (also called a **domain**) in $\mathbb{R}^2$. We denote the **boundary** of $\mathcal{D}$ by $\partial\mathcal{D}$. If $\partial\mathcal{D}$ is a collection of simple, closed curves, then each of these curves inherit a **boundary orientation** from $\mathcal{D}$. Namely, $\partial\mathcal{D}$ is oriented in such a way that, whenever the tangent line to $\partial\mathcal{D}$ is defined, the preferred unit tangent vector $\mathbf{T}$ has $\mathcal{D}$ on its left. Said another way, we choose $\mathbf{T}$ so that rotating $\mathbf{T}$ counterclockwise by $\pi/2$ radians gives an inward-pointing normal vector. Under these assumptions, we let $\mathbf{N}$ denote the **outward-pointing unit normal vector** — obtained from $\mathbf{T}$ by a clockwise rotation of $\pi/2$ radians.

Remark. Consider the region $\mathcal{D}$ in Figure 5.6. The boundary $\partial\mathcal{D}$ consists of three simple closed curves. The outer curve is oriented counterclockwise, so that $\mathcal{D}$ stays on the left side of the curve; the inner boundary components are oriented clockwise for the same reason.

With this vocabulary established, we can state Green's theorem.

Theorem 5.10: Green's theorem

Let $\mathcal{D}$ be a closed, bounded domain in $\mathbb{R}^2$ whose boundary $\partial\mathcal{D}$ is a collection of simple closed curves, appropriately oriented. Then

$$\oint_{\partial\mathcal{D}} \mathbf{F}\cdot d\mathbf{r} = \iint_{\mathcal{D}} \operatorname{curl}_z(\mathbf{F})\, dA,$$

where $\mathbf{F}$ is a vector field defined on some superset of $\mathcal{D}$.

By recalling the definition of $\operatorname{curl}_z(\mathbf{F})$ and using our alternative notation for vector line integrals, we can also write the conclusion of Green's theorem as

$$\boxed{\oint_{\partial\mathcal{D}} P\,dx + Q\,dy = \iint_{\mathcal{D}} \left(\frac{\partial Q}{\partial x} - \frac{\partial P}{\partial y}\right) dA}. \tag{5.5}$$

Idea of proof of Theorem 5.10. Our statement of Green's theorem doesn't place very many restrictions on the domain $\mathcal{D}$, which makes a careful proof of the statement somewhat beyond our scope. However, we can outline the key points.

First, we claim that any region $\mathcal{D}$ satisfying the hypotheses of Green's theorem may be decomposed into *simply connected* subregions. That is, we may subdivide $\mathcal{D}$ into subregions, each of which has only one

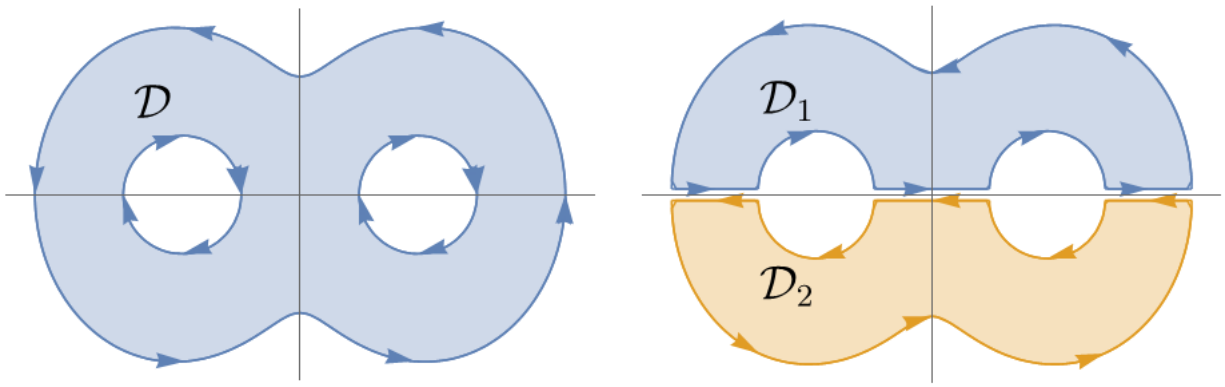

Figure 5.6: By partitioning $\mathcal{D}$ into $\mathcal{D}_1$ and $\mathcal{D}_2$, we may write both $\oint_{\partial\mathcal{D}} \mathbf{F}\cdot d\mathbf{r}$ and $\iint_{\mathcal{D}} \operatorname{curl}_z(\mathbf{F})\, dA$ as sums of integrals.

boundary component. For instance, Figure 5.6 shows a region $\mathcal{D}$ being decomposed into regions $\mathcal{D}_1$ and $\mathcal{D}_2$. By the additivity of multiple integration we know that

$$\iint_{\mathcal{D}} \operatorname{curl}_z(\mathbf{F})\, dA = \iint_{\mathcal{D}_1} \operatorname{curl}_z(\mathbf{F})\, dA + \iint_{\mathcal{D}_2} \operatorname{curl}_z(\mathbf{F})\, dA.$$

On the other hand, the orientations of our boundaries ensure that

$$\oint_{\mathcal{D}} \mathbf{F}\cdot d\mathbf{r} = \oint_{\mathcal{D}_1} \mathbf{F}\cdot d\mathbf{r} + \oint_{\mathcal{D}_2} \mathbf{F}\cdot d\mathbf{r}.$$

In particular, we see three horizontal line segments in each of $\partial\mathcal{D}_1$ and $\partial\mathcal{D}_2$ which are not part of $\partial\mathcal{D}$. Integrating along these segments will contribute to $\oint_{\mathcal{D}_1} \mathbf{F}\cdot d\mathbf{r}$ and $\oint_{\mathcal{D}_2} \mathbf{F}\cdot d\mathbf{r}$, but these contributions will cancel each other out, since $\partial\mathcal{D}_1$ and $\partial\mathcal{D}_2$ are oppositely oriented along these segments.

So it is sufficient to prove Green's theorem over simply connected regions. Our next simplification is similar in spirit, but should elicit much more suspicion. We claim that, up to some sort of limiting argument, it's sufficient to prove Green's theorem over rectangles. Unfortunately, the limiting argument is too tedious for us to include here. Loosely, instead of directly showing that

$$\oint_{\partial\mathcal{D}} \mathbf{F}\cdot d\mathbf{r} = \iint_{\mathcal{D}} \operatorname{curl}_z(\mathbf{F})\, dA,$$

we would show that, for every $\epsilon > 0$,

$$\left| \oint_{\partial\mathcal{D}} \mathbf{F}\cdot d\mathbf{r} - \iint_{\mathcal{D}} \operatorname{curl}_z(\mathbf{F})\, dA \right| < \epsilon.$$

To do this, we would select rectangles $R_1, \ldots, R_N$ contained in $\mathcal{D}$ such that

$$\left| \left(\sum_{k=1}^{N} \oint_{\partial R_k} \mathbf{F}\cdot d\mathbf{r} \right) - \oint_{\partial\mathcal{D}} \mathbf{F}\cdot d\mathbf{r} \right| < \frac{\epsilon}{2} \quad \text{and} \quad \left| \left(\sum_{k=1}^{N} \iint_{R_k} \operatorname{curl}_z(\mathbf{F})\, ds \right) - \iint_{\mathcal{D}} \operatorname{curl}_z(\mathbf{F})\, ds \right| < \frac{\epsilon}{2}.$$

Finally, we would apply Green's theorem to each rectangle. However, we'll punt on the details of this argument[2].

Exercise 5.11. Fill in the details of the above limiting argument. (This might be more complicated than I'm letting on; I don't think I've ever actually written this argument out in full detail.)

[2]You might complain that we could have done this rectangle argument without assuming that $\mathcal{D}$ is simply connected. This is a valid complaint. We just wrote simple connectivity out as a separate step so that we could get comfortable with the idea of partitioning $\mathcal{D}$ before thinking about a *sequence* of partitions, as required by the rectangles.

At long last, we prove Green's theorem under the assumption that $\mathcal{D} = [a,b] \times [c,d]$ is a rectangle in $\mathbb{R}^2$. Under this assumption, the result follows from a simple application of the Fundamental Theorem of Calculus. First, we divide $\partial\mathcal{D}$ into four line segments:

$$\mathcal{C}_1 = \{b\} \times [c,d], \quad \mathcal{C}_2 = [a,b] \times \{d\}, \quad \mathcal{C}_3 = \{a\} \times [c,d], \quad \mathcal{C}_4 = [a,b] \times \{c\}.$$

We have parametrizations for these four segments as follows:

Segment	$\mathbf{r}(t)$	Domain for t	$\mathbf{r}'(t)$	$\mathbf{F}(\mathbf{r}(t)) \cdot \mathbf{r}'(t)$
$\mathcal{C}_1$	$\mathbf{r}_1(t) = (b,t)$	$c \le t \le d$	$\mathbf{r}_1'(t) = \langle 0,1 \rangle$	$F_2(b,t)$
$\mathcal{C}_2$	$\mathbf{r}_2(t) = (-t,d)$	$-b \le t \le -a$	$\mathbf{r}_2'(t) = \langle -1,0 \rangle$	$-F_1(-t,d)$
$\mathcal{C}_3$	$\mathbf{r}_3(t) = (a,-t)$	$-d \le t \le -c$	$\mathbf{r}_3'(t) = \langle 0,-1 \rangle$	$-F_2(a,-t)$
$\mathcal{C}_4$	$\mathbf{r}_4(t) = (t,c)$	$a \le t \le b$	$\mathbf{r}_4'(t) = \langle 1,0 \rangle$	$F_1(t,c)$

Notice that $\mathcal{C}_2$ and $\mathcal{C}_3$ are parametrized "backwards," to match the counterclockwise orientation of $\partial\mathcal{D}$. Thus we find that

$$\oint_{\partial\mathcal{D}} \mathbf{F} \cdot d\mathbf{r} = \sum_{k=1}^{4} \int_{\mathcal{C}_k} \mathbf{F} \cdot d\mathbf{r} = \int_c^d F_2(b,t)\,dt - \int_{-b}^{-a} F_1(-t,d)\,dt - \int_{-d}^{-c} F_2(a,-t)\,dt + \int_a^b F_1(t,c)\,dt.$$

By applying the substitution $u = -t$ to the middle two integrals, we find that

$$\oint_{\partial\mathcal{D}} \mathbf{F} \cdot d\mathbf{r} = \int_c^d (F_2(b,t) - F_2(a,t))\,dt - \int_a^b (F_1(d,t) - F_1(c,t))\,dt.$$

Now we can quickly finish the proof. Using the Fundamental Theorem of Integral Calculus, we recognize the two integrands on the right hand side of the above equation as themselves being definite integrals:

$$\begin{aligned}
\oint_{\partial\mathcal{D}} \mathbf{F} \cdot d\mathbf{r} &= \int_c^d \left(\int_a^b \frac{\partial F_2}{\partial x}(x,y)\,dx \right) dy - \int_a^b \left(\frac{\partial F_1}{\partial y}(x,y)\,dy \right) dx \\
&= \iint_{\mathcal{D}} \left(\frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y} \right) dA \\
&= \iint_{\mathcal{D}} \operatorname{curl}_z(\mathbf{F})\,dA.
\end{aligned}$$

This is precisely the equality we hoped to establish. □

Example 5.12. Let $\mathcal{C}$ be the piecewise linear path $ABCA$ in $\mathbb{R}^2$, where $A = (0,0)$, $B = (2,0)$, and $C = (2,2)$. We can use Green's theorem to compute $\int_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r}$, where $\mathbf{F}(x,y) = \langle \sin x, x^2 y^3 \rangle$. First of all, $\mathcal{C} = \partial\mathcal{D}$, where $\mathcal{D}$ is the triangle with vertices A, B, and C. Green's theorem thus tells us that

$$\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = \iint_{\mathcal{D}} \operatorname{curl}_z(\mathbf{F})\,dA = \iint_{\mathcal{D}} 2xy^3\,dA.$$

Now $\mathcal{D}$ is vertically simple (in addition to being horizontally simple), and we may write

$$\mathcal{D} = \{(x,y) : 0 \le x \le 2, 0 \le y \le x\}.$$

So

$$\begin{aligned}
\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} &= \int_0^2 \int_0^x 2xy^3\,dy\,dx = \int_0^2 \left[\frac{1}{2} x y^4 \right]_0^x dx \\
&= \int_0^2 \frac{1}{2} x^5\,dx = \left[\frac{1}{12} x^6 \right]_0^2 = \boxed{\frac{16}{3}}.
\end{aligned}$$

This is a much quicker process than writing three separate line integrals, which would have been required to evaluate $\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r}$ directly.

Example 5.13. A famous application of Green's theorem allows us to compute the area of a domain $\mathcal{D}$ in $\mathbb{R}^2$ with a line integral. We compute the area of a domain $\mathcal{D}$ via the formula $\operatorname{area}(\mathcal{D}) = \iint_{\mathcal{D}} 1\,dA$, so in order

to use Green's theorem we'd like to find a vector field $\mathbf{F}\colon \mathcal{D} \to \mathbb{R}^2$ with $\text{curl}_z(\mathbf{F}) = 1$. Note that this can be accomplished via

$$\mathbf{F}(x,y) = \langle -\lambda y, (1-\lambda)x \rangle,$$

since we then have

$$\text{curl}_z(\mathbf{F}) = \frac{\partial}{\partial x}((1-\lambda)x) - \frac{\partial}{\partial y}(-\lambda y) = (1-\lambda) - (-\lambda) = 1.$$

In particular, by taking $\lambda = 0$ we find that

$$\boxed{\text{area}(\mathcal{D}) = \oint_{\partial \mathcal{D}} x \, dy.}$$

So, for instance, we may compute the area of the ellipse $\mathcal{D} = \left\{(x,y) : \left(\frac{x}{a}\right)^2 + \left(\frac{y}{b}\right)^2 \le 1\right\}$ by evaluating a line integral along $\partial\mathcal{D}$:

$$\begin{aligned}\text{area}(\mathcal{D}) &= \oint_{\partial \mathcal{D}} x \, dy = \oint_{\partial \mathcal{D}} \langle 0, x \rangle \cdot d\mathbf{r} = \int_0^{2\pi} \langle 0, a\cos t \rangle \cdot \langle -a\sin t, b\cos t \rangle \, dt \\ &= \int_0^{2\pi} ab\cos^2 t \, dt = \frac{ab}{2}\int_0^{2\pi} (1+\cos(2t)) \, dt = \pi ab,\end{aligned}$$

where we use the fact that the integral of $\cos(2t)$ over $[0, 2\pi]$ vanishes.

Green's theorem is an FTIC. While most of its ink was spent on making simplifying assumptions, the linchpin in the proof of Green's theorem was the single-variable Fundamental Theorem of Integral Calculus, which says that

$$F(b) - F(a) = \int_a^b F'(x) \, dx.$$

Notice that on the left we evaluate a function, but only on the boundary of some domain. Moving to the right side of the equation, we differentiate the function, and then compute the integral of this derivative across the entire domain. We can think of this as *trading a boundary for a derivative*. This phenomenon is also present in Green's theorem. The formula is

$$\oint_{\partial \mathcal{D}} \mathbf{F} \cdot d\mathbf{r} = \iint_{\mathcal{D}} \text{curl}_z(\mathbf{F}) \, dA,$$

in which we see a function being integrated over the boundary of a domain on the left, and on the right we see that function's derivative being integrated over the entire domain. This perspective — that boundaries can be traded for derivatives — motivates much of the remainder of our course.

5.3.2 The flux version of Green's theorem

Green's theorem is a beautiful generalization of the Fundamental Theorem of Integral Calculus to two dimensions, but it may leave you with a question. We said that FTICs should relate an integral over a boundary to an integral of a derivative over a region. Green's theorem gives an FTIC where the derivative in question is $\text{curl}_z(\mathbf{F})$, but this is not the only derivative which a vector field on $\mathbb{R}^2$ admits. In this subsection we want to think about a version of Green's theorem where the derivative we use is $\text{div}\,\mathbf{F}$. The justification for this result is completely analogous to that for Green's theorem, so we simply state the result and invite you to prove it.

Theorem 5.14: The flux version of Green's theorem

Let $\mathcal{D}$ be a domain in $\mathbb{R}^2$ whose boundary $\partial\mathcal{D}$ is a collection of simple closed curves, appropriately oriented. Then

$$\oint_{\partial \mathcal{D}} (\mathbf{F} \cdot \mathbf{N}) \, ds = \iint_{\mathcal{D}} \text{div}(\mathbf{F}) \, dA,$$

where $\mathbf{F}$ is a vector field defined on some superset of $\mathcal{D}$.

Exercise 5.15. Prove Theorem 5.14, at least in the case where $\mathcal{D}$ is a rectangle.

Example 5.16. In Example 5.9 we computed the flux of the vector fields

$$\mathbf{F}_1 = \langle 0, 1\rangle, \quad \mathbf{F}_2 = \langle x, -y\rangle, \quad \text{and} \quad \mathbf{F}_3 = \langle x, y\rangle.$$

through the curve $\mathcal{C}$ determined by $x^2 + 4y^2 = 16$. By realizing $\mathcal{C}$ as the boundary of the region $\mathcal{D}$ determined by $x^2 + 4y^2 \leq 16$, we can use Theorem 5.14 to compute these values again. We first note that

$$\operatorname{div}(\mathbf{F}_1) = 0, \quad \operatorname{div}(\mathbf{F}_2) = 1 + (-1) = 0, \quad \text{and} \quad \operatorname{div}(\mathbf{F}_3) = 1 + 1 = 2.$$

So we immediately see that

$$\oint_{\mathcal{C}} (\mathbf{F}_1 \cdot \mathbf{N})\, ds = \iint_{\mathcal{D}} 0\, dA = 0 \quad \text{and} \quad \oint_{\mathcal{C}} (\mathbf{F}_2 \cdot \mathbf{N})\, ds = \iint_{\mathcal{D}} 0\, dA = 0,$$

while

$$\oint_{\mathcal{C}} (\mathbf{F}_3 \cdot \mathbf{N})\, ds = \iint_{\mathcal{D}} 2\, dA = 2\operatorname{area}(\mathcal{D}).$$

Now we may rewrite the defining inequality of $\mathcal{D}$ as $(\frac{x}{4})^2 + (\frac{y}{2})^2 \leq 1$, and then Example 5.13 tells us that $\operatorname{area}(\mathcal{D}) = 8\pi$. So the flux of $\mathbb{F}_3$ through $\mathcal{C}$ is 16π, matching our computation in Example 5.9.

5.3.3 The geometric meaning of curl and divergence

Finally, Theorems 5.10 and 5.14 allow us to loosely justify the names of the derivatives $\operatorname{curl}_z(\mathbf{F})$ and $\operatorname{div}(\mathbf{F})$. We will give a geometric explanation for the divergence, and leave the explanation of curl as an exercise.

Given a vector field $\mathbf{F}\colon \mathbb{R}^2 \to \mathbb{R}^2$ and a point $(x_0, y_0) \in \mathbb{R}^2$, a reasonable way to estimate the extent to which $\mathbf{F}$ is diverging at (x_0, y_0) is as follows: draw a small disc

$$\mathcal{D}_r = \{(x, y) : (x - x_0)^2 + (y - y_0)^2 \leq r^2\}$$

centered on (x_0, y_0) and ask for the total flux of $\mathbf{F}$ through $\partial\mathcal{D}_r$. This tells us how much $\mathbf{F}$ is running away from (x_0, y_0). Now as we let r tend to 0, so that the disc becomes arbitrarily small, the flux of $\mathbf{F}$ through $\partial\mathcal{D}_r$ will of course also tend to 0. But we can ask for some sort of *average* flux by dividing this total flux by the area of $\mathcal{D}_r$:

$$\overline{\text{flux}}_r = \frac{1}{\operatorname{area}(\mathcal{D}_r)} \oint_{\partial\mathcal{D}_r} (\mathbf{F} \cdot \mathbf{N})\, ds.$$

Our informal claim is that the divergence of $\mathbf{F}$ at a point (x_0, y_0) should be some sort of "instantaneous average flux" at (x_0, y_0). That is, $\operatorname{div}(\mathbf{F})(x_0, y_0)$ should be given by the limit of $\overline{\text{flux}}_r$ as $r \to 0$. But then, using Theorem 5.14, we find that

$$\lim_{r\to 0} \overline{\text{flux}}_r = \lim_{r\to 0} \left(\frac{1}{\operatorname{area}(\mathcal{D}_r)} \oint_{\partial\mathcal{D}_r} (\mathbf{F} \cdot \mathbf{N})\, ds \right) = \lim_{r\to 0} \left(\frac{1}{\operatorname{area}(\mathcal{D}_r)} \iint_{\mathcal{D}_r} \operatorname{div}(\mathbf{F})\, dA \right) = \operatorname{div}(\mathbf{F})(x_0, y_0). \tag{5.6}$$

So we can think of the divergence of $\mathbf{F}$ at (x_0, y_0) as measuring the average flux of $\mathbf{F}$ through the boundary of $\mathcal{D}$, where $\mathcal{D}$ is some arbitrarily small region around (x_0, y_0).

Remark. Perhaps the most suspicious part of Equation 5.6 is the final equality. In words, this says that, as we let a region shrink down to a point, the average value of a function over that region will tend towards the function's value at the point. The corresponding equality in one variable is

$$\lim_{r\to 0} \left(\frac{1}{2r} \int_{x_0 - r}^{x_0 + r} f(x)\, dx \right) = f(x_0).$$

I encourage you to prove this using the usual Fundamental Theorem of Integral Calculus, under the assumption that f is continuous.

Exercise 5.17. Carry out a similar argument, using Green's theorem, to convince yourself that evaluating $\operatorname{curl}_z(\mathbf{F})$ at (x_0, y_0) gives the average counterclockwise rotation of $\mathbf{F}$ along the boundary of some arbitrarily small region about (x_0, y_0).

6 Week 6

This week we'll tell a story for surfaces which is very similar to the story told last week for curves. Specifically, we will define what it means to integrate a function over a surface, and then what it means to integrate a vector field. On Friday, the miracle of Green's theorem is repeated (and enhanced) by Stokes' theorem, which establishes a fundamental relationship between surface integrals and line integrals.

6.1 Day 15: Surface integrals of functions

Goals

By the end of today's class, we should be able to do the following.

1. Parametrize various basic surfaces in $\mathbb{R}^3$.
2. Justify the definition of **surface integrals** of scalar functions over parametrized surfaces in $\mathbb{R}^3$.
3. Use surface integrals to compute **surface areas** and **masses**.

The key idea behind line integrals of scalar functions is that a curve $\mathcal{C}$, while it may live in any space $\mathbb{R}^n$, is fundamentally a one-dimensional object, and thus we may treat the integral of a function over $\mathcal{C}$ as a single-variable integral. By analogous reasoning, if $\mathcal{S} \subset \mathbb{R}^n$ is a **surface** — loosely meaning that $\mathcal{S}$ appears to be two-dimensional — then we should have some sort of "double integral" for functions defined on $\mathcal{S}$. Today we'll try to make sense of this for parametrized surfaces.

6.1.1 Parametrized surfaces

Definition. A **simple parametrized surface in** $\mathbb{R}^n$, $n \geq 2$, is a piecewise C^1-mapping

$$G : \mathcal{D} \to \mathbb{R}^n,$$

where $\mathcal{D} \subset \mathbb{R}^2$ is a region[1] in $\mathbb{R}^2$ (often — but not always — assumed to be closed and bounded) and G is injective when restricted to the interior $\mathring{\mathcal{D}}$ of $\mathcal{D}$. We say that G **is a parametrization of** $\mathcal{S}$ if $G(\mathcal{D}) = \mathcal{S}$. We say that G is **regular** if the $n \times 2$ Jacobian matrix

$$G_*(p) := \begin{pmatrix} \frac{\partial G_1}{\partial u}(p) & \frac{\partial G_1}{\partial v}(p) \\ \vdots & \vdots \\ \frac{\partial G_n}{\partial u}(p) & \frac{\partial G_n}{\partial v}(p) \end{pmatrix}$$

has full rank, for all points $p \in \mathring{\mathcal{D}}$.

Remark. If you took 32AH or have otherwise seen enough linear algebra, then the notion of rank for a matrix should be familiar. In case it's not, we point out that $G_*(p)$ has full rank if and only if the column vectors[2]

$$\frac{\partial G}{\partial u} = \left\langle \frac{\partial G_1}{\partial u}, \cdots, \frac{\partial G_n}{\partial u} \right\rangle \quad \text{and} \quad \frac{\partial G}{\partial v} = \left\langle \frac{\partial G_1}{\partial v}, \cdots, \frac{\partial G_n}{\partial v} \right\rangle$$

are linearly independent. In case $n = 3$, this is equivalent to the condition that $\frac{\partial G}{\partial u} \times \frac{\partial G}{\partial v}$ is a non-vanishing vector field on the interior of $\mathcal{D}$.

Continuing to mimic the story we told for curves, we can define **reparametrizations** of parametrized surfaces. However, rather than adapt the definition of reparametrizations which we used for curves, let's observe that reparametrizing a curve amounts to choosing a coordinate system $\phi : J \to I$ on the domain of

[1] We haven't given a careful definition of the word *region* — and we won't — but we assume that regions have *nonempty interiors*. For a region $\mathcal{D}$ in $\mathbb{R}^2$, this just means that it's possible to draw a disc of nonzero area somewhere in $\mathcal{D}$.

[2] We haven't ever specified whether we think of vectors as columns or rows, mostly because this doesn't seem like a very illuminating discussion. But we have tried to be consistent about using angle brackets for row vectors and parentheses for column vectors. Maybe it's confusing that we said column vectors here and then wrote two row vectors. We just mean that $\frac{\partial G}{\partial u}$ and $\frac{\partial G}{\partial v}$ are the columns of G_*, but we write them as rows for the sake of space.

our curve. By analogy, suppose $G\colon \mathcal{D} \to \mathbb{R}^n$ is a parametrized surface, and that $\Phi\colon R \to \mathbb{R}^2$ is a coordinate system on $\mathcal{D}$, for some rectangle $R \subset \mathbb{R}^2$. Then

$$G \circ \Phi\colon \Phi^{-1}(\mathcal{D}) \to \mathbb{R}^n$$

is a parametrized surface whose image agrees with that of G, so we say that we have a reparametrization of G. As with reparametrizations of curves, we say that Φ is an **orientation-preserving** reparametrization if $\det \Phi_* > 0$, and call Φ **orientation-reversing** if $\det \Phi_* < 0$. (In the case of curves, note that the Jacobian determinant of a coordinate system $\phi : J \to I$ is given by $\det \phi_* = \phi'$, so this matches what we wrote last week.)

Remark. Hopefully the domain we're using for $G \circ \Phi$ isn't too much of a stumbling block. The main idea here is that while the image of Φ contains $\mathcal{D}$, it may contain other points as well. In general, $(G \circ \Phi)(p)$ will only be defined for a point $p \in R$ if $\Phi(p) \in \mathcal{D}$, and thus the domain of $G \circ \Phi$ is $\Phi^{-1}(\mathcal{D}) \subset R$.

Example 6.1. Using our knowledge of the cylindrical coordinate system, parametrizing a cylinder centered on the z-axis in $\mathbb{R}^3$ is no problem at all. Indeed, we can parametrize the cylinder $\mathcal{C}_R = \{(x,y,z) : x^2 + y^2 = R^2\}$ with the map $G\colon [0,2\pi]_\theta \times \mathbb{R}_z \to \mathbb{R}^3$ defined by

$$G(\theta, z) := (R\cos\theta, R\sin\theta, z).$$

When restricted to $(0,2\pi)_\theta \times \mathbb{R}_z$, this map is injective, and we notice that the Jacobian matrix

$$G_* = \begin{pmatrix} -R\sin\theta & 0 \\ R\cos\theta & 0 \\ 0 & 1 \end{pmatrix}$$

has full rank, for all choices of θ and z. Indeed, the cross product

$$\frac{\partial G}{\partial \theta} \times \frac{\partial G}{\partial z} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ -R\sin\theta & R\cos\theta & 0 \\ 0 & 0 & 1 \end{vmatrix} = \langle R\cos\theta, R\sin\theta, 0\rangle$$

never vanishes, and thus G is a regular parametrization. We point out that, while G is inspired by cylindrical coordinates, this is a fundamentally different object — it's a parametrization of a surface.

Example 6.2. Continuing to draw motivation from familiar coordinate systems, we can easily parametrize spheres in $\mathbb{R}^3$. Consider the sphere $\mathcal{S}_R = \{(x,y,z) : x^2 + y^2 + z^2 = R^2\}$, for some $R > 0$. We parametrize this sphere with the map $G\colon [0,2\pi]_\theta \times [0,\pi]_\varphi \to \mathbb{R}^3$ defined by

$$G(\theta, \varphi) = (R\sin\varphi\cos\theta, R\sin\varphi\sin\theta, R\cos\varphi).$$

Note that G is injective when restricted to $(0,2\pi) \times (0,\pi)$, and that

$$\begin{aligned} \frac{\partial G}{\partial \theta} \times \frac{\partial G}{\partial \varphi} &= \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ -R\sin\varphi\sin\theta & R\sin\varphi\cos\theta & 0 \\ R\cos\varphi\cos\theta & R\cos\varphi\sin\theta & -R\sin\varphi \end{vmatrix} \\ &= \left\langle -R^2\sin^2\varphi\cos\theta, -R^2\sin^2\varphi\sin\theta, -R^2\sin\varphi\cos\varphi \right\rangle. \end{aligned}$$

It may not be obvious that this cross product never vanishes, but we can verify this by computing the vector's magnitude:

$$\begin{aligned} \left\| \frac{\partial G}{\partial \theta} \times \frac{\partial G}{\partial \varphi} \right\|^2 &= (-R^2\sin^2\varphi\cos\theta)^2 + (-R^2\sin^2\varphi\sin\theta)^2 + (-R^2\sin\varphi\cos\varphi)^2 \\ &= R^4\sin^4\varphi + R^4\sin^2\varphi\cos^2\varphi = R^4\sin^2\varphi. \end{aligned}$$

We have assumed that $R \neq 0$, and we have $0 < \varphi < \pi$ on the interior of our domain, so $\sin\varphi \neq 0$. We conclude that the cross product does not vanish on the interior of our domain. From this it follows that the Jacobian matrix of G has full rank, and thus G is regular. Again, this parametrization is inspired by a familiar coordinate system, but is not itself a coordinate system.

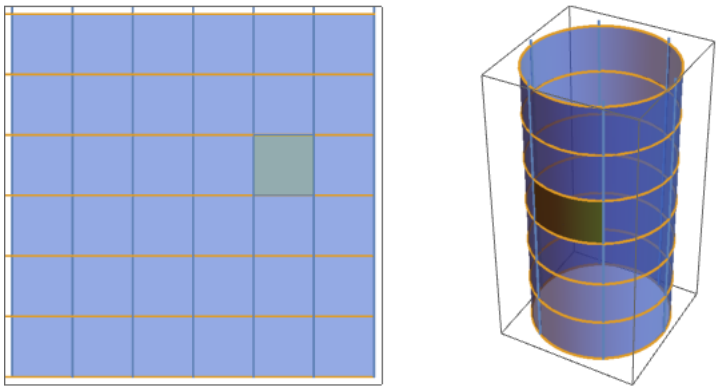

Figure 6.1: We can partition a parametrized surface $\mathcal{S}$ by partitioning the domain $\mathcal{D}$ of a parametrization $G\colon \mathcal{D} \to \mathbb{R}^3$.

Example 6.3. Another surface in $\mathbb{R}^3$ which is easily parametrized is given by the graph of a differentiable function $f : \mathcal{D} \subset \mathbb{R}^2 \to \mathbb{R}$. Since this graph is defined by the equation $z = f(x, y)$, we simply define our parametrization by

$$G(u, v) = (u, v, f(u, v)),$$

for all $(u, v) \in \mathcal{D}$. Notice that G is injective on all of $\mathcal{D}$ (not just on the interior). The Jacobian matrix of G is given by

$$G_* = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ \frac{\partial f}{\partial u} & \frac{\partial f}{\partial v} \end{pmatrix},$$

verifying that G is a C^1-mapping. So G is a simple parametrized surface. To make sure that G is regular, we compute

$$\frac{\partial G}{\partial u} \times \frac{\partial G}{\partial v} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ 1 & 0 & \frac{\partial f}{\partial u} \\ 0 & 1 & \frac{\partial f}{\partial v} \end{vmatrix} = \left\langle -\frac{\partial f}{\partial u}, -\frac{\partial f}{\partial v}, 1 \right\rangle.$$

Notice that this cross product will never vanish, regardless of the partial derivatives f_u and f_v. So G is a regular parametrization.

6.1.2 Integrating scalar functions

As stated at the outset, what we really want to do is define some notion of "double integral" for functions $f : \mathcal{S} \to \mathbb{R}$ defined on parametrized surfaces in $\mathbb{R}^3$. We can do this in a manner very similar to what we did for line integrals. Namely, if we have a regular parametrization $G\colon \mathcal{D} \to \mathbb{R}^3$ of $\mathcal{S}$, then we can partition $\mathcal{S}$ by partitioning $\mathcal{D}$ and then using G to push this partition over to $\mathcal{S}$. See Figure 6.1.

Once we have a partition, we could estimate the value of our function f over each pseudobox in the partition by plugging in some test point from the pseudobox, or by taking the infimum of all values of f over the pseudobox, or by taking the supremum of all such values, or whatever. But, as was the case with line integrals, we then find that we don't know the area of the pseudobox, so we have to estimate it by taking a planar approximation which agrees with our surface at the four corners of the pseudobox. As our partition becomes arbitrarily fine, this approximation should become arbitrarily good.

But all of this sounds like a lot of tedious work, and, given our history with integration, we can tell where the story is going. When we came up with Theorem 5.2 we noticed that this formula fits into an analogy with the change of variables theorem. Let's extend that analogy. We have

$$\begin{array}{ccc} \mathbb{R}^n \xrightarrow{\Phi} \mathbb{R}^n & \qquad I \xrightarrow{\mathbf{r}} \mathbb{R}^n & \qquad \mathcal{D} \xrightarrow{G} \mathbb{R}^3 \\ \quad\downarrow f & \qquad\quad \downarrow f & \qquad\quad \downarrow f \\ \quad\mathbb{R} & \qquad\quad \mathbb{R} & \qquad\quad \mathbb{R} \end{array} \tag{6.1}$$

The first two diagrams depict the situation for the change of variables theorem and line integrals, respectively. In both cases, we use our coordinate system or parametrization to "pull back" our integral, instead integrating the function $f \circ \Phi$ (or $f \circ \mathbf{r}$) over a more reasonable domain. In the process of pulling back, we pick up a distortion factor. By the same reasoning, we should have a formula that looks like

$$\iint_{\mathcal{S}} f \, dS = \iint_{\mathcal{D}} (f \circ G) \circledast \, dA,$$

where $\circledast$ is some to-be-determined distortion factor. So it remains to determine the distortion factor $\circledast$.

To sort out what $\circledast$ should be, we think back to the two-variable version of the change-of-variables theorem. We found that the coordinate system Φ turned a box with side lengths Δu and Δv into a pseudobox which could be approximated by a parallelogram. Specifically, we approximated the pseudobox with the parallelogram spanned by $\Delta u \frac{\partial \Phi}{\partial u}$ and $\Delta v \frac{\partial \Phi}{\partial v}$, and then discovered that this parallelogram has area

$$\left| \det \left(\Delta u \frac{\partial \Phi}{\partial u} \quad \Delta v \frac{\partial \Phi}{\partial v} \right) \right| = \Delta u \Delta v \left| \det \left(\frac{\partial \Phi}{\partial u} \quad \frac{\partial \Phi}{\partial v} \right) \right| = \Delta u \Delta v |\det \Phi_*|.$$

This is how we ended up with $|\det \Phi_*|$ as our fudge factor for change of variables. By very similar reasoning, a regular parametrization G will map a box with side lengths Δu, Δv to a pseudobox which we approximate with the parallelogram spanned by $\Delta u \frac{\partial G}{\partial u}$ and $\Delta v \frac{\partial G}{\partial v}$. As before, we approximate the area of the pseudobox by computing the area of this parallelogram. The hangup this time is that $\frac{\partial G}{\partial u}$ and $\frac{\partial G}{\partial v}$ are vectors in $\mathbb{R}^3$, so we can't take the determinant of the matrix with these vectors as its columns. Instead, we remember that the area of the parallelogram spanned by two vectors in $\mathbb{R}^3$ is given by the magnitude of the cross product of those vectors. So the area of our pseudobox is approximately

$$\left\| \left(\Delta u \frac{\partial G}{\partial u} \right) \times \left(\Delta v \frac{\partial G}{\partial v} \right) \right\| = \Delta u \Delta v \left\| \frac{\partial G}{\partial u} \times \frac{\partial G}{\partial v} \right\|.$$

We conclude that, if $G : \mathcal{D} \to \mathbb{R}^3$ is a regular parametrization of a surface, then the area distortion factor is $\|G_u \times G_v\|$. This leads to the following definition.

Definition. Let $G : \mathcal{D} \to \mathbb{R}^3$ be a regular, simple parametrization of a surface $\mathcal{S}$. For any function $f : \mathcal{S} \to \mathbb{R}$, we define the surface integral of f over $\mathcal{S}$ to be

$$\iint_{\mathcal{S}} f \, dS := \iint_{\mathcal{D}} (f \circ G) \left\| \frac{\partial G}{\partial u} \times \frac{\partial G}{\partial v} \right\| \, dA, \tag{6.2}$$

provided the integral on the right exists.

Exercise 6.4. Show that surface integrals are well-defined. Specifically, if $\mathcal{S}$ is a regular surface, then show that the value of equation 6.2 is independent of what parametrization of $\mathcal{S}$ we use.

Remark. This definition is philosophically different from all of our previous integral definitions. We didn't want to deal with all the details of partitions and lower and upper sums and whatnot, so we just defined our integral according to an equation which really ought to be true. Exercise 6.4 justifies this approach.

As was the case for line integrals, our formula for computing surface integrals assumes that we have a regular parametrization, but we often want to integrate over objects (in this case, surfaces) which don't admit a regular parametrization. We won't develop this theory carefully, but we'll see through some examples that we can do the natural thing: we treat surface integration as additive. For instance, to compute a surface integral over the surface of a cube, we compute six regular surface integrals — one for each face of the cube. Note that the regular pieces into which we break our surface are allowed to have nontrivial intersection, but this intersection must have volume zero. Again, we won't be super careful about this.

Remark. We have only defined scalar surface integrals for surfaces in $\mathbb{R}^3$ (unlike scalar line integrals, which we defined for curves in any dimension). Had we taken an approach which involved partitions, we could have written down a definition which works for surfaces in higher dimensions, but writing down a formula like (6.2) would have been difficult, since the cross product only makes sense in dimension 3.

6.1.3 Examples

Example 6.5. We can use a surface integral to compute surface area. If $\mathcal{S}$ has finite surface area, then this area *ought* to be calculated by

$$\text{area}(\mathcal{S}) = \iint_{\mathcal{S}} 1 \, dS.$$

Let's take $\mathcal{S}$ to be that portion of the cone $z^2 = x^2 + y^2$ which lies between the planes $z = 1$ and $z = 4$. This is pretty easily parametrized by modifying cylindrical coordinates. Specifically, we can define $G : [1,4]_r \times [0,2\pi]_\theta \to \mathbb{R}^3$ by

$$G(r,\theta) = (r\cos\theta, r\sin\theta, r).$$

We'll leave it as an exercise to check that G is injective when restricted to its interior, but this is certainly believable. We now compute the distortion factor for this parametrization. First,

$$\frac{\partial G}{\partial r} \times \frac{\partial G}{\partial \theta} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \cos\theta & \sin\theta & 1 \\ -r\sin\theta & r\cos\theta & 0 \end{vmatrix} = \langle -r\cos\theta, -r\sin\theta, r \rangle .$$

Then

$$\left\| \frac{\partial G}{\partial r} \times \frac{\partial G}{\partial \theta} \right\|^2 = 2r^2,$$

so

$$\begin{aligned} \text{area}(\mathcal{S}) = \iint_{\mathcal{S}} 1 \, dS &= \iint_{\mathcal{D}} \left\| \frac{\partial G}{\partial r} \times \frac{\partial G}{\partial \theta} \right\| dA \\ &= \int_1^4 \int_0^{2\pi} \sqrt{2} r \, d\theta \, dr = \int_0^{2\pi} \left[\frac{\sqrt{2}}{2} r^2 \right]_1^4 d\theta \\ &= \int_0^{2\pi} \frac{\sqrt{2}}{2}(16-1) \, d\theta = \boxed{15\sqrt{2}\pi}. \end{aligned}$$

Example 6.6. The equation $y = 9 - z^2$ determines a parabolic cylinder in $\mathbb{R}^3$. Let's consider the portion $\mathcal{S}$ of this cylinder which satisfies $0 \le x \le 3$ and $0 \le z \le 3$. This surface can be parametrized by

$$G(u,v) = (u, 9 - v^2, v)$$

over the domain $[0,3]_u \times [0,3]_v$. Note that if $G(u_1,v_1) = G(u_2,v_2)$ then, by investigating the first and third components, we see that $(u_1,v_1) = (u_2,v_2)$. So G is injective. We also see that

$$\frac{\partial G}{\partial u} \times \frac{\partial G}{\partial v} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ 1 & 0 & 0 \\ 0 & -2v & 1 \end{vmatrix} = \langle 0, -1, -2v \rangle .$$

Because the third component of this vector field is -1, we see that $G_u \times G_v$ is nowhere-vanishing. So G is a regular parametrization[3]. Now consider the function $f(x,y,z) = z$. We can restrict this to a function $f : \mathcal{S} \to \mathbb{R}$ and we find that the composition $f \circ G : [0,3] \times [0,3] \to \mathbb{R}$ is given by

$$(f \circ G)(u,v) = f(u, 9 - v^2, v) = v.$$

So

$$\begin{aligned} \iint_{\mathcal{S}} f \, dS &= \iint_{[0,3]\times[0,3]} (f \circ G) \| \langle 0, -1, -2v \rangle \| \, dA = \int_0^3 \int_0^3 v\sqrt{1+4v^2} \, dv \, du \\ &= 3 \int_0^3 v\sqrt{1+4v^2} \, dv = \left[\frac{3}{12}(1+4v^2)^{3/2} \right]_0^3 = \boxed{\frac{1}{4}(37\sqrt{37} - 1)}. \end{aligned}$$

Gross.

[3] Indeed, since y is a function of x and z, our earlier discussion about parametrizing graphs of functions applies.

Example 6.7. As a teaser for our next section, let's start with a vector field and a surface and write down a scalar surface integral. For now we won't say why we care about this particular scalar surface integral, but we'll answer that question next time. Consider the vector field $\mathbf{F} = \langle x, y, z\rangle$ and the sphere $\mathcal{S}_R = \{(x, y, z) : x^2 + y^2 + z^2 = R^2\}$, for some $R > 0$. The unit vector

$$\mathbf{N} = \frac{1}{\sqrt{x^2 + y^2 + z^2}} \langle x, y, z\rangle$$

is everywhere perpendicular to $\mathcal{S}_R$. Let's compute the surface integral of the (scalar-valued) function $\mathbf{F} \cdot \mathbf{N}$ over $\mathcal{S}_R$. First, we need to compute $\mathbf{F} \cdot \mathbf{N}$:

$$f(x, y, z) := \mathbf{F}(x, y, z) \cdot \mathbf{N}(x, y, z) = \frac{1}{\sqrt{x^2 + y^2 + z^2}} \langle x, y, z\rangle \cdot \langle x, y, z\rangle = \frac{x^2 + y^2 + z^2}{\sqrt{x^2 + y^2 + z^2}} = \sqrt{x^2 + y^2 + z^2}.$$

Now what we really care about is the restriction of f to $\mathcal{S}_R$. But on $\mathcal{S}_R$ we know that $x^2 + y^2 + z^2 = R^2$, so $f|_{\mathcal{S}_R} \equiv R$. So we're really just computing the integral of the constant function R. Experience tells us that the result should just be R times the surface area of $\mathcal{S}_R$, which gives $4\pi R^3$, but let's do this the hard way. As explained above, we have a parametrization $G : [0, 2\pi] \times [0, \pi] \to \mathbb{R}^3$, defined by

$$G(\theta, \varphi) = (R \sin\varphi \cos\theta, R \sin\varphi \sin\theta, R \cos\varphi),$$

and we find that $\|G_\theta \times G_\varphi\| = R^2 \sin\varphi$. So

$$\begin{aligned}\iint_{\mathcal{S}_R} (\mathbf{F} \cdot \mathbf{N})\, dS &= \iint_{[0,2\pi]\times[0,\pi]} (R)(R^2 \sin\varphi)\, dA = \int_0^{2\pi} \int_0^{\pi} R^3 \sin\varphi \; d\varphi \; d\theta \\ &= \int_0^{2\pi} R^3 [-\cos\varphi]_0^{\pi} \; d\theta = \int_0^{2\pi} R^3 2 \; d\theta = \boxed{4\pi R^3},\end{aligned}$$

as expected.

6.2 Day 16: Surface integrals of vector fields

Goals

By the end of today's class, we should be able to do the following.

1. Justify the definition of **surface integrals** of vector fields over parametrized surfaces.
2. Use vector surface integrals to measure the **flux** of a vector field through a surface.

Continuing to follow last week's outline, we now move from scalar surface integrals to vector surface integrals. If we don't look too closely, the stories have a lot of similarities: we have a vector field $\mathbf{F}$ which is defined all around the object over which we want to integrate (last week this object was a curve, and now it's a surface). To pull off the integration, we write down a scalar function which is determined by the vector field and the object (last week it was $\mathbf{F} \cdot \mathbf{T}$, and this week we'll see that it's $\mathbf{F} \cdot \mathbf{N}$).

But it's worth pointing out where this analogy break down. Loosely, given a vector field $\mathbf{F}$ on $\mathbb{R}^n$, we can integrate $\mathbf{F}$ over objects which have dimension 1, or those which have *codimension* 1. Objects which have dimension 1 are just curves[4], and indeed we can define the line integral of $\mathbf{F}$ over a curve in any dimension n. In this case, the integral of $\mathbf{F}$ measures the extent to which $\mathbf{F}$ pushes along the object. On the other hand, an object has codimension 1 if it has one fewer dimension than the ambient space. In $\mathbb{R}^2$, curves have codimension 1; when we compute the work done by $\mathbf{F}$ along $\mathcal{C}$, we're thinking of $\mathcal{C}$ as having dimension 1, and when we compute the flux of $\mathbf{F}$ through $\mathcal{C}$ we're thinking of $\mathcal{C}$ as having codimension 1. In $\mathbb{R}^3$, surfaces have codimension 1, and we call objects in $\mathbb{R}^n$, $n > 3$, with codimension 1 *hypersurfaces*. For an object $\mathcal{S} \subset \mathbb{R}^n$ with codimension 1, we can ask how much $\mathbf{F}$ pushes *through* $\mathcal{S}$ rather than asking how much $\mathbf{F}$ pushes *along* $\mathcal{S}$.

Our story today takes place in $\mathbb{R}^3$. We already know how to compute the integral of a vector field along a curve in $\mathbb{R}^3$, so we'll focus on computing the integral of a vector field over a surface. That is, we already know the dimension 1 setup, so we focus on the *co*dimension 1 situation. According to the above discussion, this should measure the flux of $\mathbf{F}$ through $\mathcal{S}$.

[4]We haven't said what we mean by the *dimension* of a subset of $\mathbb{R}^n$, so this is not at all a careful statement. We're just trying to outline the story.

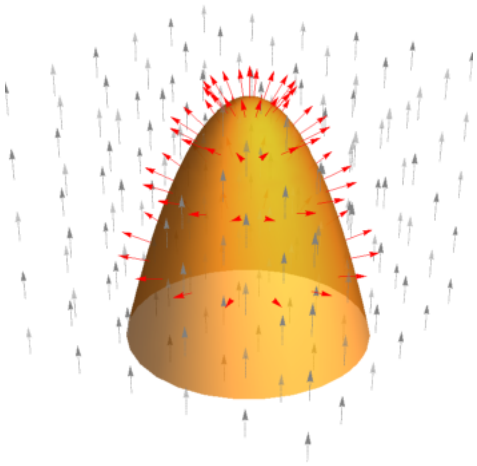

Figure 6.2: The surface integral $\iint_{\mathcal{S}} \mathbf{F} \cdot d\mathbf{S}$ measures the total flux of $\mathbf{F}$ through $\mathcal{S}$.

6.2.1 Vector surface integrals

Following the above description, we are just about ready to give the definition of the surface integral of a vector field. But first we point out that, as was the case for vector line integrals, the vector surface integral requires the surface over which we're integrating to be **oriented**. Namely, we will assume that a preferred unit normal vector $\mathbf{N}$ has been chosen at all points of $\mathcal{S}$, and in fact assume that this choice is continuous[5]. Not all surfaces admit such a choice — for instance, we can't reasonably choose a normal vector along the edges of a cube in $\mathbb{R}^3$ — but, unlike with curves, this can happen for interesting reasons. A very famous example is the Möbius strip, a surface which permits a normal vector at each point, but for which we cannot make a continuous global choice of normal vector.

Assuming that $\mathcal{S}$ is oriented, we observe that the extent to which a vector field $\mathbf{F}$ pushes through $\mathcal{S}$ at a point $p \in \mathcal{S}$ is measured by $\mathbf{F}(p) \cdot \mathbf{N}(p)$. This motivates the following definition.

Definition. Let $\mathcal{S} \subset \mathbb{R}^3$ be an oriented surface which admits a regular, simple parametrization, and let $\mathbf{F} \colon A \to \mathbb{R}^3$ be a vector field defined on some superset A of $\mathcal{S}$. The **surface integral of $\mathbf{F}$ along $\mathcal{S}$** is defined by

$$\iint_{\mathcal{S}} \mathbf{F} \cdot d\mathbf{S} := \iint_{\mathcal{S}} (\mathbf{F} \cdot \mathbf{N}) \, dS, \tag{6.3}$$

where $\mathbf{N}$ is the oriented unit normal vector to $\mathcal{S}$, provided the integral on the right exists. The quantity $\iint_{\mathcal{S}} \mathbf{F} \cdot d\mathbf{S}$ is the **flux** of $\mathbf{F}$ through $\mathcal{S}$.

Now suppose we have a regular, simple parametrization $G \colon \mathcal{D} \to \mathbb{R}^3$ of $\mathcal{S}$. We claim that the vector field $G_u \times G_v$ is perpendicular to $\mathcal{S}$. Indeed, the tangent plane to $\mathcal{S}$ at $G(u,v) \in \mathcal{S}$ is spanned by G_u and G_v, and the cross product $G_u \times G_v$ is perpendicular to each of these vectors. We will say that G is an **oriented parametrization** of $\mathcal{S}$ if $G_u \times G_v$ points in the same direction as $\mathbf{N}$ throughout the interior of $\mathcal{D}$. In this case we have

$$\mathbf{N}(G(u,v)) = \frac{(G_u \times G_v)(u,v)}{\|(G_u \times G_v)(u,v)\|}$$

for all $(u,v) \in \mathcal{D}$. Given a regular, simple parametrization G, let's introduce the notation

$$\boxed{\mathbf{n}(u,v) := \frac{\partial G}{\partial u} \times \frac{\partial G}{\partial v},}$$

to save ourselves the pain of writing the cross product so many times. We make some observations about this notation:

[5] It's probably worth thinking about what this means; we won't.

1. There's not really an established convention for how this notation should work. So in some texts the lowercase **n** refers to the unit normal vector which we denote by **N**, while others agree with what we do here. Be careful out there.
2. Unlike **N**, the vector $\mathbf{n}(u,v)$ depends on our parametrization G. In particular, if G is a negatively-oriented parametrization, then **n** won't even point in the same direction as **N**.
3. In general, if a regular parametrization is given for a surface, but no orientation is otherwise specified for $\mathcal{S}$, then we will assume that G is positively oriented. In this case, $\mathbf{N} = \mathbf{n}/\|\mathbf{n}\|$.

By combining Equations 6.2 and 6.3, we conclude that

$$\boxed{\iint_{\mathcal{S}} \mathbf{F}\cdot d\mathbf{S} = \iint_{\mathcal{D}} \left(\mathbf{F}(G(u,v))\cdot \frac{\mathbf{n}(u,v)}{\|\mathbf{n}(u,v)\|}\right)\|\mathbf{n}(u,v)\|\ dA = \iint_{\mathcal{D}} \mathbf{F}(G(u,v))\cdot \mathbf{n}(u,v)\ dA}. \tag{6.4}$$

So, just like line integrals, vector surface integrals are somewhat easier to compute than their scalar counterparts, since we don't have a vector magnitude to consider.

6.2.2 Examples

Example 6.8. Figure 6.2 depicts the surface

$$\mathcal{S} = \{(x,y,z) : z = 4 - x^2 - y^2, z \geq 0\}$$

and the vector field $\mathbf{F} = \langle 0,0,1\rangle$. Let's compute the flux of **F** through $\mathcal{S}$, where $\mathcal{S}$ is oriented with an upward-pointing normal vector. To do this, consider the parametrization

$$G(r,\theta) := (r\cos\theta, r\sin\theta, 4-r^2),$$

defined on the domain $\mathcal{D} = [0,2]_r \times [0,2\pi]_\theta$. Notice that this parametrization is inspired by cylindrical coordinates, since in that coordinate system the defining equation becomes $z = 4-r^2$ and the inequality $z \geq 0$ implies $r \leq 2$. One may check that G is injective on $(0,2)\times(0,2\pi)$, and we see that

$$\mathbf{n}(r,\theta) = \frac{\partial G}{\partial r}\times\frac{\partial G}{\partial \theta} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \cos\theta & \sin\theta & -2r \\ -r\sin\theta & r\cos\theta & 0 \end{vmatrix} = \langle 2r^2\cos\theta, 2r^2\sin\theta, r\rangle.$$

Then $\|\mathbf{n}(r,\theta)\| = r\sqrt{4r^2+1}$ does not vanish on $(0,2)\times(0,2\pi)$, so G is a regular parametrization. In fact, we could get away without the magnitude computation by noticing that the final component of $\mathbf{n}(r,\theta)$ cannot vanish when $r \neq 0$. Moreover, since $r > 0$ on the interior of our domain, $\mathbf{n}(r,\theta)$ is upward-pointing, and thus our parametrization is positively-oriented. So we find that

$$\begin{aligned}\iint_{\mathcal{S}} \mathbf{F}\cdot d\mathbf{S} &= \iint_{\mathcal{D}} \mathbf{F}(G(r,\theta))\cdot\mathbf{n}(r,\theta)\ dr\ d\theta = \int_0^{2\pi}\int_0^2 \langle 0,0,1\rangle\cdot\langle 2r^2\cos\theta, 2r^2\sin\theta, r\rangle\ dr\ d\theta \\ &= \int_0^{2\pi}\int_0^2 r\ dr\ d\theta = \boxed{4\pi}.\end{aligned}$$

The final equality follows from the fact that the given iterated integral computes the double integral of 1 over a disc of radius 2. But this will just give the area of the disc, which is 4π.

Example 6.9. Let's think about the ellipsoid

$$\mathcal{S} = \left\{(x,y,z) : \left(\frac{x}{4}\right)^2 + \left(\frac{y}{3}\right)^2 + \left(\frac{z}{2}\right)^2 = 1\right\}.$$

We can compute the flux of the vector field $\mathbf{F} = z\mathbf{i}$ through $\mathcal{S}$ with an outward-pointing normal vector. By modifying spherical coordinates, we find that

$$G(\theta,\varphi) = (4\sin\varphi\cos\theta, 3\sin\varphi\sin\theta, 2\cos\varphi)$$

is a parametrization of $\mathcal{S}$, with domain $\mathcal{D} = [0,2\pi]_\theta \times [0,\pi]_\varphi$. The injectivity of G on the interior of its domain follows from the usual considerations for spherical coordinates, so we check that G is regular. We have

$$\mathbf{n}(\theta,\varphi) = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ -4\sin\varphi\sin\theta & 3\sin\varphi\cos\theta & 0 \\ 4\cos\varphi\cos\theta & 3\cos\varphi\sin\theta & -2\sin\varphi \end{vmatrix} = \langle -6\sin^2\varphi\cos\theta, -8\sin^2\varphi\sin\theta, -12\sin\varphi\cos\varphi\rangle.$$

Computing the magnitude of the vector field doesn't seem fun, so let's make some observations. On $(0,2\pi)_\theta \times (0,\pi)_\varphi$, the function $\sin\varphi$ does not vanish, and thus if $\mathbf{n}(\theta,\varphi) = \mathbf{0}$, we must have $\cos\varphi = 0$, according to the final component. But if $\cos\varphi = 0$, then $\sin^2\varphi = 1$, meaning that the first two components vanish only if $\cos\theta = 0$ and $\sin\theta = 0$. Because these equations cannot hold simultaneously, we conclude that $\mathbf{n}(\theta,\varphi)$ is nonzero on the interior of our domain. So G is a regular parametrization of $\mathcal{S}$.

But there's a catch! By comparing $\mathbf{n}(\theta,\varphi)$ with $G(\theta,\varphi)$, we notice that G has an inward-pointing normal vector. So G is a negatively-oriented parametrization of $\mathcal{S}$. All this means for us from a practical standpoint is that we pick up a negative sign on our flux computation:

$$\begin{aligned}\iint_{\mathcal{S}} \mathbf{F}\cdot d\mathbf{S} &= -\iint_{\mathcal{D}} \mathbf{F}(G(u,v))\cdot\mathbf{n}(\theta,\varphi)\,dA \\ &= -\int_0^{2\pi}\int_0^{\pi} \langle 2\cos\varphi,0,0\rangle\cdot\langle -6\sin^2\varphi\cos\theta,-8\sin^2\varphi\sin\theta,-12\sin\varphi\cos\varphi\rangle \\ &= -\int_0^{2\pi}\int_0^{\pi} -12\cos\varphi\sin^2\varphi\cos\theta\;d\varphi\;d\theta = \int_0^{2\pi}\left[4\cos\theta\sin^3\varphi\right]_0^{\pi}\;d\theta = \boxed{0}.\end{aligned}$$

So the net flux of $\mathbf{F}$ through $\mathcal{S}$ is 0. You should try to convince yourself that this is a reasonable conclusion.

6.3 Day 17: Stokes' theorem

Goals

By the end of today's class, we should be able to do the following.

1. Prove **Stokes' theorem** for oriented, parametrized surfaces in $\mathbb{R}^3$.
2. Use Stokes' theorem to pass back and forth between line integrals and surface integrals.
3. Give a geometric interpretation of the curl of a vector field.

Last week's classes culminated in Green's theorem, which is a version of the Fundamental Theorem of Integral Calculus for regions in $\mathbb{R}^2$. With the understanding of surfaces and surface integrals that we've gained this week, we can realize a double integral over a region in $\mathbb{R}^2$ as a surface integral in $\mathbb{R}^3$. This begs the question: does Green's theorem actually hold for *any* vector surface integral in $\mathbb{R}^3$? That is, if $\mathcal{C}\subset\mathbb{R}^3$ is the oriented boundary of some surface $\mathcal{S}\subset\mathbb{R}^3$, can we still trade boundaries for derivatives?

Today we'll show that the answer is yes; our generalization of Green's theorem is **Stokes' theorem**.

6.3.1 Stokes' theorem

Before stating our result, we need to extend our definition of **boundary orientation** just a bit. Specifically, we previously defined the boundary orientation given by a region $\mathcal{D}\subset\mathbb{R}^2$ to the curves which make up its boundary; we extend this notion to surfaces in $\mathbb{R}^3$.

Definition. Let $\mathcal{S}\subset\mathbb{R}^3$ be an oriented surface which admits a simple, regular parametrization, and let $\mathbf{N}$ be the preferred unit normal vector on $\mathcal{S}$. Assume that the boundary $\partial\mathcal{S}$ is made up of curves which admit simple, regular parametrizations. The **boundary orientation** on these curves prefers the unit tangent vector $\mathbf{T}$ with the (somewhat informal) property that $\mathbf{T}\times\mathbf{N}$ points out of $\mathcal{S}$.

Remark. Here are some comments on the boundary orientation:

1. The definition given here extends the definition for regions in the plane. Namely, regions in the plane have preferred unit normal vector $\mathbf{k}$ (as surfaces in $\mathbb{R}^3$), so $\mathbf{T}\times\mathbf{k}$ points out of the region whenever we choose $\mathbf{T}$ so that the region is on the left.
2. Another sense in which our definition extends the previous one is as follows: if we traverse $\partial\mathcal{S}$ with the correct orientation, always standing on the same side of $\mathcal{S}$ as the normal vector $\mathbf{N}$, then $\mathcal{S}$ will be to our left. This is much more easily explained with pictures, so we'll draw some in class.
3. In case you've seen the *Frenet frame*, we point out that $\mathbf{N}$ is the unit normal vector to $\mathcal{S}$, not to $\mathcal{C}$, so $\mathbf{T}\times\mathbf{N}$ need not be parallel to the binormal vector $\mathbf{B}$.

4. To make this notion of boundary orientation work in higher dimensions, we'll end up using some algebraic machinery to be discussed later. But the basic idea will be the same: the orientation on the boundary is chosen so that it pairs well with an outward-pointing vector for the original object.

Theorem 6.10: Stokes' theorem in $\mathbb{R}^3$

Let $\mathcal{S} \subset \mathbb{R}^3$ be an oriented surface which admits a regular, simple parametrization, and let $\mathbf{F}: A \to \mathbb{R}^3$ be a C^1-vector field on some open superset A of $\mathcal{S}$. Then

$$\oint_{\partial \mathcal{S}} \mathbf{F} \cdot d\mathbf{r} = \iint_{\mathcal{S}} \operatorname{curl}(\mathbf{F}) \cdot d\mathbf{S}, \tag{6.5}$$

where $\partial \mathcal{S}$ carries the boundary orientation.

Remark. If $\partial \mathcal{S} = \emptyset$ — for instance, if $\mathcal{S}$ is a sphere — then the integral on the left of (6.5) is defined to be 0.

Proof idea for Theorem 6.10. The basic idea of our proof is that we may use a sufficiently nice parametrization $G: \mathcal{D} \to \mathbb{R}^3$ of $\mathcal{S}$ to rewrite Equation 6.5 as an equality of integrals in the plane. Once in the plane, we can simply apply Green's theorem. The details of this argument are a little technical, but we'll give some of them.

Let's suppose we have a parametrization $G: \mathcal{D} \to \mathbb{R}^3$ which is injective on all of $\mathcal{D}$, and for which the Jacobian matrix G_* has full rank on all of $\mathcal{D}$. Essentially we are asking that, in addition to being a regular, simple parametrization, G is regular and simple along $\partial \mathcal{D}$ as well. We point out that not all of the surfaces to which Theorem 6.10 applies admit such a parametrization (e.g., a sphere admits no such parametrization). We are just using this as a simplifying assumption. In another simplifying assumption, we suppose that the boundary $\partial \mathcal{D}$ is a single curve which admits a regular, simple parametrization.

With such a G chosen, we define a vector field $\mathbf{X}: \mathcal{D} \to \mathbb{R}^2$ by

$$\mathbf{X}(u, v) := \left\langle \mathbf{F}(G(u, v)) \cdot \frac{\partial G}{\partial u}, \mathbf{F}(G(u, v)) \cdot \frac{\partial G}{\partial v} \right\rangle.$$

This choice of $\mathbf{X}$ is sorely lacking in motivation; don't worry if it seems like this came out of nowhere — it sort of did! Our plan is for $\mathbf{X}$ to be some kind of avatar for $\mathbf{F}$ on the uv-plane.

Now choose a regular, simple parametrization $\mathbf{r}: I \to \mathbb{R}^2_{u,v}$ of $\partial \mathcal{D}$, and write $\mathbf{r}(t) = (u(t), v(t))$. Then $(G \circ \mathbf{r}): I \to \mathbb{R}^3$ is a parametrization of $\partial \mathcal{S}$, and from the chain rule it follows that

$$(G \circ \mathbf{r})'(t) = G_*(\mathbf{r}(t))\, \mathbf{r}'(t) = u'(t) \frac{\partial G}{\partial u}(\mathbf{r}(t)) + v'(t) \frac{\partial G}{\partial v}(\mathbf{r}(t)).$$

Because G is regular, the vectors G_u and G_v are linearly independent, and thus this derivative can only vanish if the scalar quantities $u'(t)$ and $v'(t)$ simultaneously vanish. But this doesn't happen, because we assumed that $\mathbf{r}(t)$ is regular. So $G \circ \mathbf{r}$ is a regular parametrization of $\partial \mathcal{S}$, and we can use it to compute the integral of $\mathbf{F}$ along $\partial \mathcal{S}$:

$$\begin{aligned}
\oint_{\partial \mathcal{S}} \mathbf{F} \cdot d\mathbf{r} &= \int_I \mathbf{F}((G \circ \mathbf{r})(t)) \cdot (G \circ \mathbf{r})'(t)\, dt \\
&= \int_I \mathbf{F}(G(\mathbf{r}(t))) \cdot \left(u'(t) \frac{\partial G}{\partial u} + v'(t) \frac{\partial G}{\partial v} \right) dt \\
&= \int_I \left(u'(t)\, \mathbf{F}(G(\mathbf{r}(t))) \cdot \frac{\partial G}{\partial u} + v'(t)\, \mathbf{F}(G(\mathbf{r}(t))) \cdot \frac{\partial G}{\partial v} \right) dt \\
&= \int_I \left\langle \mathbf{F}(G(\mathbf{r}(t))) \cdot \frac{\partial G}{\partial u}, \mathbf{F}(G(\mathbf{r}(t))) \cdot \frac{\partial G}{\partial v} \right\rangle \cdot \left\langle u'(t), v'(t) \right\rangle dt
\end{aligned}$$

This last expression might look kind of scary, but by comparing with the definition of $\mathbf{X}(u, v)$, we see that we have

$$\oint_{\partial \mathcal{S}} \mathbf{F} \cdot d\mathbf{r} = \int_I \mathbf{X}(\mathbf{r}(t)) \cdot \mathbf{r}'(t)\, dt = \oint_{\partial \mathcal{D}} \mathbf{X} \cdot d\mathbf{r} = \iint_{\mathcal{D}} \operatorname{curl}_z(\mathbf{X})\, dA, \tag{6.6}$$

where the last step follows from Green's theorem in the plane. So we've transformed the line integral $\oint_{\partial S} \mathbf{F}\cdot d\mathbf{r}$ into a double integral of a curl. This is a good first step, and now we need to realize our double integral as a surface integral. To this end, we notice that $\text{curl}_z(\mathbf{X})$ may be computed as

$$\begin{aligned}\text{curl}_z(\mathbf{X}) &= \frac{\partial}{\partial u}\left(\mathbf{F}(G(u,v))\cdot\frac{\partial G}{\partial v}\right) - \frac{\partial}{\partial v}\left(\mathbf{F}(G(u,v))\cdot\frac{\partial G}{\partial u}\right)\\ &= \left(\mathbf{F}_*\frac{\partial G}{\partial u}\right)\cdot\frac{\partial G}{\partial v} + \left(\mathbf{F}(G(u,v))\cdot\frac{\partial^2 G}{\partial u\partial v}\right) - \left(\mathbf{F}_*\frac{\partial G}{\partial v}\right)\cdot\frac{\partial G}{\partial u} - \left(\mathbf{F}(G(u,v))\cdot\frac{\partial^2 G}{\partial v\partial u}\right)\\ &= \left(\mathbf{F}_*\frac{\partial G}{\partial u}\right)\cdot\frac{\partial G}{\partial v} - \left(\mathbf{F}_*\frac{\partial G}{\partial v}\right)\cdot\frac{\partial G}{\partial u},\end{aligned}$$

where

$$\mathbf{F}_* = \begin{pmatrix} \frac{\partial F_1}{\partial x} & \frac{\partial F_1}{\partial y} & \frac{\partial F_1}{\partial z} \\ \frac{\partial F_2}{\partial x} & \frac{\partial F_2}{\partial y} & \frac{\partial F_2}{\partial z} \\ \frac{\partial F_3}{\partial x} & \frac{\partial F_3}{\partial y} & \frac{\partial F_3}{\partial z} \end{pmatrix}$$

is the Jacobian matrix of the map $\mathbf{F}\colon A\to\mathbb{R}^3$. Further unwinding of $\text{curl}_z(\mathbf{X})$ is somewhat tedious, but you are invited to verify the following fact.

Exercise 6.11. Using the definition of $\mathbf{F}_*$, expand the above computation to show that

$$\text{curl}_z(\mathbf{X}) = \text{curl}(\mathbf{F})\cdot\left(\frac{\partial G}{\partial u}\times\frac{\partial G}{\partial v}\right).$$

This is great, because substitution into (6.6) tells us that

$$\oint_{\partial\mathcal{S}} \mathbf{F}\cdot d\mathbf{r} = \iint_{\mathcal{D}} \text{curl}(\mathbf{F})(G(u,v))\cdot\mathbf{n}(u,v)\,dA.$$

But the integral on the right is, by definition, the surface integral $\iint_S \text{curl}(\mathbf{F})\cdot d\mathbf{S}$. So we've verified (6.5) (under our simplifying assumptions), and our proof is concluded. □

6.3.2 Examples

Example 6.12. Stokes' theorem can be particularly effective when we're asked to compute a line integral over a curve which is presented as a boundary. Consider the piecewise-linear curve $\mathcal{C}$ which is the boundary of $\mathcal{T}$, the triangle in $\mathbb{R}^3$ with vertices at $A=(3,0,0)$, $B=(0,2,0)$, and $C=(0,0,1)$. Say we want to compute the integral of

$$\mathbf{F}(x,y,z) := \left\langle \sin(x^2), e^{y^2}+x^2, z^4+2x^2\right\rangle$$

over $\mathcal{C}$, with $\mathcal{C}$ oriented as $ABCA$. We'll do this using Stokes' theorem. First, notice that

$$\text{curl}(\mathbf{F}) = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \partial/\partial x & \partial/\partial y & \partial/\partial z \\ \sin(x^2) & e^{y^2}+x^2 & z^4+2x^2 \end{vmatrix} = \left\langle 0, -\frac{\partial}{\partial x}(z^4+2x^2), \frac{\partial}{\partial x}(e^{y^2}+x^2)\right\rangle = \langle 0,-4x,2x\rangle.$$

So the vector field we need to integrate has simplified. We now have

$$\oint_{\mathcal{C}} \mathbf{F}\cdot d\mathbf{r} = \iint_{\mathcal{T}} \text{curl}(\mathbf{F})\cdot d\mathbf{S} = \iint_{\mathcal{T}} (\text{curl}(\mathbf{F})\cdot\mathbf{N})\,dS,$$

where $\mathbf{N}$ is the unit normal vector to $\mathcal{T}$. Now you may recall that, in general, the plane passing through the points $(a,0,0)$, $(0,b,0)$, and $(0,0,c)$, with none of a,b,c equal to 0, has equation

$$bc\,x + ac\,y + ab\,z = abc.$$

So $\mathcal{T}$ is contained in the plane $2x+3y+6z=6$. Having written this equation in standard form, we quickly notice that $\langle 2,3,6\rangle$ is a normal vector for $\mathcal{T}$ (though not a unit normal vector). But we have

$$\text{curl}(\mathbf{F})\cdot\langle 2,3,6\rangle = (0)(2)+(-4x)(3)+(2x)(6) = 0,$$

so we must have $\text{curl}(\mathbf{F}) \cdot \mathbf{N} = 0$. We conclude that

$$\boxed{\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = \iint_{\mathcal{T}} (\text{curl}(\mathbf{F}) \cdot \mathbf{N})\, dS = 0.}$$

Notice that not only did we avoid computing the line integral of $\mathbf{F}$ directly, we were also able to avoid computing a surface integral of $\text{curl}(\mathbf{F})$, since this vector field turned out to be tangent to our surface. Also notice that, since the flux of $\text{curl}(\mathbf{F})$ through $\mathcal{T}$ is 0, the orientation of $\mathcal{T}$ didn't matter.

Example 6.13. Here's an example you probably hoped to never see again. In our two most recent homework assignments we considered the surface

$$\mathcal{S}_{h,R} := \{(x, y, f(x, y)) : (x, y) \in \mathcal{D}_R\},$$

where $\mathcal{D}_R$ is the disc of radius R in $\mathbb{R}^2$ centered at the origin, and

$$f(x, y) := h - \frac{h}{R^2}(x^2 + y^2).$$

Specifically, we thought about the flux of the vector field

$$\mathbf{F}(x, y, z) := \langle 2y - 1, x(y - 2), 1 - xz \rangle$$

through $\mathcal{S}_{h,R}$. We now know that this flux can be computed as a vector surface integral:

$$\text{flux} = \iint_{\mathcal{S}_{h,R}} \mathbf{F} \cdot d\mathbf{S}.$$

Now in the most recent homework assignment, you identified a vector potential $\mathbf{G}$ for $\mathbf{F}$, and this affords us an opportunity to use Stokes' theorem, as we may write

$$\iint_{\mathcal{S}_{h,R}} \mathbf{F} \cdot d\mathbf{S} = \iint_{\mathcal{S}_{h,R}} \text{curl}(\mathbf{G}) \cdot d\mathbf{S} = \oint_{\partial \mathcal{S}_{h,R}} \mathbf{G} \cdot d\mathbf{r}.$$

At this point we could recall the particular vector field $\mathbf{G}$ that you found and compute this line integral explicitly, but let's consider an alternative approach. We notice that, whatever the values of h and R, the boundary $\partial \mathcal{S}_{h,R}$ is precisely the boundary $\partial \mathcal{D}_R$ in $\mathbb{R}^2 \subset \mathbb{R}^3$. This is a fact you verified in homework 2, and holds because $f(x, y) = 0$ for any $(x, y) \in \partial \mathcal{D}_R$. Treating $\mathcal{D}_R$ as a surface with upward-pointing normal vector (so that the boundary orientation induced by $\mathcal{D}_R$ on $\partial \mathcal{D}_R$ agrees with that induced by $\mathcal{S}_{h,R}$), we can use Stokes' theorem again to see that

$$\oint_{\partial \mathcal{S}_{h,r}} \mathbf{G} \cdot d\mathbf{r} = \oint_{\partial \mathcal{D}_R} \mathbf{G} \cdot d\mathbf{r} = \iint_{\mathcal{D}_R} \text{curl}(\mathbf{G}) \cdot d\mathbf{S} = \iint_{\mathcal{D}_R} \mathbf{F} \cdot d\mathbf{S}.$$

So we notice that $\mathbf{F}$ is *surface-independent*: integrating $\mathbf{F}$ over any surface with oriented boundary $\partial \mathcal{S}_{h,R}$ yields the same value as integrating $\mathbf{F}$ over any other such surface. It is now straightforward to compute the integral of $\mathbf{F}$ over $\mathcal{D}_R$. Notice that, since $z = 0$ along $\mathcal{D}_R$, $\mathbf{F}|_{\mathcal{D}_R} = \langle 2y - 1, x(y - 2), 1 \rangle$. The unit normal vector to $\mathcal{D}_R$ is given by $\mathbf{N} = \mathbf{k}$, so we find that

$$\mathbf{F} \cdot \mathbf{N} = \langle 2y - 1, x(y - 2), 1 \rangle \cdot \langle 0, 0, 1 \rangle = 1$$

along $\mathcal{D}_R$. So

$$\iint_{\mathcal{D}_R} \mathbf{F} \cdot d\mathbf{S} = \iint_{\mathcal{D}_R} (\mathbf{F} \cdot \mathbf{N})\, dS = \iint_{\mathcal{D}_R} 1\, dS = \text{area}(\mathcal{D}_R).$$

Since $\mathcal{D}_R$ is a disc with radius R, this last value is simply πR^2. So we find that

$$\boxed{\iint_{\mathcal{S}_{h,R}} \mathbf{F} \cdot d\mathbf{S} = \pi R^2,}$$

independent of the value of h.

It's worth pointing out that the computation in the above example never made use of the particular vector potential **G** that you computed; we simply needed to know that such a vector field existed. Poincaré's lemma is a particularly useful tool for these sorts of computations, as it allows us (under certain hypotheses) to conclude that a vector potential exists without needing to compute it.

Example 6.14. In the previous example we observed that integrating our vector field **F** over $\mathcal{D}_R$ would yield the same result as integrating over $\mathcal{S}_{h,R}$, since these two surfaces have the same oriented boundary. Indeed, Stokes' theorem tells us that any vector field which admits a vector potential enjoys this surface-independence. Say $\mathbf{F}\colon A \to \mathbb{R}^3$ is defined on some open subset A of $\mathbb{R}^3$ and admits a vector potential $\mathbf{G}\colon A \to \mathbb{R}^3$. Then, if $\mathcal{S}_1$ and $\mathcal{S}_2$ are a pair of surfaces contained in A, we have

$$\begin{aligned}\iint_{\mathcal{S}_1} \mathbf{F}\cdot d\mathbf{S} &= \iint_{\mathcal{S}_1} \operatorname{curl}(\mathbf{G})\cdot d\mathbf{S} = \oint_{\partial\mathcal{S}_1} \mathbf{G}\cdot d\mathbf{r} \\ &= \oint_{\partial\mathcal{S}_2} \mathbf{G}\cdot d\mathbf{r} = \iint_{\mathcal{S}_2} \operatorname{curl}(\mathbf{G})\cdot d\mathbf{S} = \iint_{\mathcal{S}_2} \mathbf{F}\cdot d\mathbf{S},\end{aligned}$$

where the second and fourth equalities follow from Stokes' theorem. So we have another necessary condition for a vector field to admit a vector potential: the field must be *surface-independent*. We will explore the sufficiency of this condition next week, as well as an analogous condition (path-independence) for conservative vector fields.

6.3.3 The geometric meaning of curl

Just as Green's theorem gave us a physical/geometric interpretation of the quantity $\operatorname{curl}_z(\mathbf{F})$, Stokes' theorem allows us to loosely understand why the vector field **F** is so named.

Remark. Sadly, I seriously doubt that we'll have time to discuss this subsection in class. But I'll leave it here and encourage you to ask about it in office hours. In any case, this is a no-pressure, just-for-fun discussion; you won't be responsible for this material on a midterm or final.

Let's start the game over, pretending that we haven't yet defined curl(**F**). Now say we have a sufficiently differentiable vector field $\mathbf{F}\colon \mathbb{R}^3 \to \mathbb{R}^3$ and we want to decide how much **F** "curls" at a point $p_0 = (x_0, y_0, z_0)$. The transition from two dimensions to three makes this question somewhat peculiar. In two dimensions, this question could be answered with a number — positive numbers indicate counterclockwise rotation, negative numbers mean clockwise rotation, and zero means no rotation at all. But in $\mathbb{R}^3$, our vector field could be rotating in a number of ways; in some sense, rotation is a phenomenon we can only easily understand in two dimensions. Whatever curl(**F**) is, we want it to permit the following interpretation: given any plane P_0 through our point p_0, curl(**F**) should be able to tell us how much **F** is curling in P_0.

So we need curl(**F**) to be an object which pairs naturally with planes; vector fields are a good candidate. Specifically, if curl(**F**) is a vector field, then for every plane P_0 which passes through p_0, we can consider the number

$$\operatorname{curl}(\mathbf{F})(p_0)\cdot\mathbf{N}_{P_0},$$

where $\mathbf{N}_{P_0}$ is an oriented unit normal vector to P_0. We can try to define curl(**F**) in such a way that this number will measure the extent to which **F** is spinning about p_0 in the plane P_0.

So we've decided what type of algebraic object we'll use to encode rotation: a vector field. Now we need to actually define the vector field curl(**F**). One way to measure the extent to which **F** rotates about p_0 in P_0 is to ask, on average, how much work **F** does along a circle of radius $R > 0$ centered on p_0 and contained in P_0. The smaller we choose $R > 0$, the more accurate this measurement of the curl should be, so let's say we want[6]

$$\operatorname{curl}(\mathbf{F})(p_0)\cdot\mathbf{N}_{P_0} = \lim_{R\to 0} \frac{1}{\operatorname{area}(\mathcal{D}_R)} \oint_{\partial\mathcal{D}_R} \mathbf{F}\cdot d\mathbf{r},$$

[6]One fishy bit about this average that we're computing is the fact that we're dividing by the area of $\mathcal{D}_R$, rather than the length of $\partial\mathcal{D}_R$. This is basically just to make the dimensions check out. The line integral over $\partial\mathcal{D}_R$ computes work, and therefore will have dimension $[\text{mass}]\cdot[\text{length}]^2/[\text{time}]^2$. Dividing by the length of $\partial\mathcal{D}_R$ leaves a factor of length in our quantity, so the quantity will vanish as R tends to 0. Dividing instead by area avoids this problem.

where $\mathcal{D}_R \subset P_0$ is the disc of radius R centered at p_0 in P_0. Now Stokes' theorem tells us that this is equal to

$$\lim_{R\to 0} \frac{1}{\text{area}(\mathcal{D}_R)} \iint_{\mathcal{D}_R} (\nabla \times \mathbf{F}) \cdot d\mathbf{S},$$

where $\nabla \times \mathbf{F}$ is just shorthand for

$$\nabla \times \mathbf{F} := \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \partial/\partial x & \partial/\partial y & \partial/\partial z \\ F_1 & F_2 & F_3 \end{vmatrix} := \left\langle \frac{\partial F_3}{\partial y} - \frac{\partial F_2}{\partial z}, \frac{\partial F_1}{\partial z} - \frac{\partial F_3}{\partial x}, \frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y} \right\rangle.$$

(Since we're starting over in our search for curl($\mathbf{F}$), we don't want to call this vector field curl right now.) Now $\mathcal{D}_R$, being contained in P_0, has unit normal vector given by $\mathbf{N}_{P_0}$, so

$$\lim_{R\to 0} \frac{1}{\text{area}(\mathcal{D}_R)} \iint_{\mathcal{D}_R} (\nabla \times \mathbf{F}) \cdot d\mathbf{S} = \lim_{R\to 0} \frac{1}{\text{area}(\mathcal{D}_R)} \iint_{\mathcal{D}_R} \left((\nabla \times \mathbf{F}) \cdot \mathbf{N}_{P_0} \right) \, dS.$$

So we're taking the limit of the average value of $(\nabla \times \mathbf{F}) \cdot \mathbf{N}_{P_0}$ over the disc $\mathcal{D}_R$ as this disc shrinks down to the point p_0. The result of this limit is just $(\nabla \times \mathbf{F})(p_0) \cdot \mathbf{N}_{P_0}$ — the function whose average value we're computing, evaluated at the point to which all the discs are shrinking.

Okay, so what does all of this mean? We said we wanted to define the vector field curl($\mathbf{F}$) in such a way that $\text{curl}(\mathbf{F})(p_0) \cdot \mathbf{N}_{P_0}$ would measure the extent to which $\mathbf{F}$ is curling around p_0 in P_0. But then we found, using Stokes' theorem, that this curling is measured by $(\nabla \times \mathbf{F})(p_0) \cdot \mathbf{N}_{P_0}$. So $\nabla \times \mathbf{F}$ is a perfectly good candidate for curl($\mathbf{F}$)! Taking this much more geometric approach, we find it natural to define $\text{curl}(\mathbf{F}) := \nabla \times \mathbf{F}$, where again the quantity on the right is just algebraic shorthand.

Remember that we initially defined curl($\mathbf{F}$) as a purely algebraic object, meant to encode a necessary condition for conservativity. The point of this subsection is that Stokes' theorem allows us to also interpret curl($\mathbf{F}$) as a measurement for how much $\mathbf{F}$ rotates at the points in its domain — a feature which is fitting for a vector field with this name.

Exercise 6.15. As an open-ended thought experiment, think about how you might extend this whole discussion to vector fields on $\mathbb{R}^4$, or even $\mathbb{R}^n$. We said above that Clairaut's theorem will give $\binom{4}{2} = 6$ partial differential equations which must be satisfied in order for a vector field to be conservative. These can be encoded as a single six-dimensional vector field. On the other hand, how many measurements do we need to take to determine the extent to which $\mathbf{F}$ is curling at a point $p_0 \in \mathbb{R}^4$ (or $\mathbb{R}^n$)? How do these measurements relate to the conditions given by Clairaut's theorem?

7 Week 7

We'll start this week with **Gauss's divergence theorem**, which gives a Fundamental Theorem of Integral Calculus in which surfaces play the boundary role, rather than the region role they played in Stokes' theorem. After some practice using the divergence theorem, we'll make a return to $\mathbb{R}^2$. Namely, we'll think about the curl-free condition that conservative vector fields must satisfy, and consider just how far this necessary condition is from being sufficient.

7.1 Day 18: The divergence theorem

Goals

By the end of today's class, we should be able to do the following.

1. Sketch the proof of the **divergence theorem** in any dimension.
2. Use the divergence theorem to calculate flux through a surface.

Today we'll prove another Fundamental Theorem of Integral Calculus, this time relating surface integrals to triple integrals. Recall that Stokes' theorem relates the surface integral of a derivative (specifically, the curl) of a vector field to a line integral along the boundary of the surface. On the other hand, a closed surface can bound a region, and then we can ask for an FTIC which treats the surface as a boundary. The result is the **divergence theorem**, which we will prove and use today.

7.1.1 The divergence theorem

The divergence theorem applies to surfaces which bound closed regions in $\mathbb{R}^3$; just like last week, we'll assume that our regions have nonempty interior, whatever that means. As should any FTIC, the divergence theorem allows us to trade this boundary in for a derivative: instead of integrating a vector field $\mathbf{F}$ over $\partial\mathcal{W}$, we integrate the divergence $\operatorname{div}(\mathbf{F})$ over all of $\mathcal{W}$.

Theorem 7.1: The divergence theorem

Let $\mathcal{W}$ be a closed region in $\mathbb{R}^3$ whose boundary is a (possibly disconnected) surface $\partial\mathcal{W}$, with outward-pointing normal vector. If $\mathbf{F}\colon A \to \mathbb{R}^3$ is a C^1 vector field on some open set A containing $\mathcal{W}$, then

$$\oiint_{\partial\mathcal{W}} \mathbf{F}\cdot d\mathbf{S} = \iiint_{\mathcal{W}} \operatorname{div}(\mathbf{F})\, dV. \tag{7.1}$$

Remark. Just as with Green's theorem, the divergence theorem requires that our vector field be defined on the whole of the region $\mathcal{W}$ in question. So we cannot apply the divergence theorem, for instance, to an inverse-square vector field if the point around which this vector field is centered is somewhere in $\mathcal{W}$.

Proof of Theorem 7.1. As we did for Green's theorem, we'll prove the divergence theorem for boxes and claim that this can be enhanced to a proof for a more general region $\mathcal{W}$. The basic idea is as before: given some $\epsilon > 0$, we find boxes $B_1, \dots, B_{N_\epsilon} \subset \mathcal{W}$ (for some very large integer $N_\epsilon \gg 0$) such that the integrals in (7.1) may be approximated by sums of integrals over the boxes:

$$\left| \oiint_{\partial\mathcal{W}} \mathbf{F}\cdot d\mathbf{S} - \sum_{k=1}^{N_\epsilon} \oiint_{\partial\mathcal{B}_k} \mathbf{F}\cdot d\mathbf{S} \right| < \frac{\epsilon}{2} \quad \text{and} \quad \left| \iiint_{\mathcal{W}} \operatorname{div}(\mathbf{F})\, dV - \sum_{k=1}^{N_\epsilon} \iiint_{B_k} \operatorname{div}(\mathbf{F})\, dV \right| < \frac{\epsilon}{2}.$$

Once we have these inequalities, (7.1) will hold for $\mathcal{W}$, provided it holds for all boxes (which we will prove). The tough part is probably showing that we can find the desired boxes $B_1, \dots, B_{N_\epsilon}$ in the first place.

Okay, so we need to prove that (7.1) holds under the assumption that $\mathcal{W}$ is a box in $\mathbb{R}^3$. Specifically, let's

write $\mathcal{W} = [x_0, x_1] \times [y_0, y_1] \times [z_0, z_1]$. Notice that if $\mathbf{F} = \langle F_1, F_2, F_3 \rangle$, then $\mathbf{F} = F_1\mathbf{i} + F_2\mathbf{j} + F_3\mathbf{k}$, and thus we have

$$\oiint_{\partial\mathcal{W}} \mathbf{F} \cdot d\mathbf{S} = \oiint_{\partial\mathcal{W}} (\mathbf{F} \cdot \mathbf{N})\, dS = \oiint_{\partial\mathcal{W}} F_1(\mathbf{i} \cdot \mathbf{N})\, dS + \oiint_{\partial\mathcal{W}} F_2(\mathbf{j} \cdot \mathbf{N})\, dS + \oiint_{\partial\mathcal{W}} F_3(\mathbf{k} \cdot \mathbf{N})\, dS. \tag{7.2}$$

We investigate the three surface integrals on the right individually. We begin with $\oiint_{\partial\mathcal{W}} F_1(\mathbf{i} \cdot \mathbf{N})\, dS$. Now $\partial\mathcal{W}$ has six faces, and we may summarize their unit normal vectors as follows:

Face	$\{x = x_0\}$	$\{x = x_1\}$	$\{y = y_0\}$	$\{y = y_1\}$	$\{z = z_0\}$	$\{z = z_1\}$
$\mathbf{N}$	$-\mathbf{i}$	$\mathbf{i}$	$-\mathbf{j}$	$\mathbf{j}$	$-\mathbf{k}$	$\mathbf{k}$

So the dot product $\mathbf{i} \cdot \mathbf{N}$ will only be nonzero along the faces $\{x = x_0\}$ and $\{x = x_1\}$, and we find that

$$\oiint_{\partial\mathcal{W}} F_1(\mathbf{i} \cdot \mathbf{N})\, dS = \iint_{\{x=x_1\}} F_1\, dS - \iint_{\{x=x_0\}} F_1\, dS.$$

Now the faces $\{x = x_0\}$ and $\{x = x_1\}$ admit parametrizations

$$G_0(y, z) = (x_0, y, z) \quad \text{and} \quad G_1(y, z) = (x_1, y, z),$$

respectively, where $y_0 \le y \le y_1$ and $z_0 \le z \le z_1$. It's easy to check that these are regular parametrizations with distortion factor 1, so, being kind of sloppy/skipping some steps, we observe that

$$\oiint_{\partial\mathcal{W}} F_1(\mathbf{i} \cdot \mathbf{N})\, dS = \int_{z_0}^{z_1} \int_{y_0}^{y_1} (F_1(x_1, y, z) - F_1(x_0, y, z))\, dy\, dz.$$

Just as we did in the proof of Green's theorem, we recognize the integrand as the result of a definite integral:

$$\oiint_{\partial\mathcal{W}} F_1(\mathbf{i} \cdot \mathbf{N})\, dS = \int_{z_0}^{z_1} \int_{y_0}^{y_1} \int_{x_0}^{x_1} \frac{\partial F_1}{\partial x}(x, y, z)\, dx\, dy\, dz = \iiint_{\mathcal{W}} \frac{\partial F_1}{\partial x}\, dV.$$

By completely analogous reasoning, we find that

$$\oiint_{\partial\mathcal{W}} F_2(\mathbf{j} \cdot \mathbf{N})\, dS = \iiint_{\mathcal{W}} \frac{\partial F_2}{\partial y}\, dV \quad \text{and} \quad \oiint_{\partial\mathcal{W}} F_3(\mathbf{k} \cdot \mathbf{N})\, dS = \iiint_{\mathcal{W}} \frac{\partial F_3}{\partial z}\, dV.$$

Substituting our findings back into (7.2), we see that

$$\oiint_{\partial\mathcal{W}} \mathbf{F} \cdot d\mathbf{S} = \iiint_{\mathcal{W}} \frac{\partial F_1}{\partial x}\, dV + \iiint_{\mathcal{W}} \frac{\partial F_2}{\partial y}\, dV + \iiint_{\mathcal{W}} \frac{\partial F_3}{\partial z}\, dV = \iiint_{\mathcal{W}} \operatorname{div}(\mathbf{F})\, dV,$$

exactly as desired. □

Just as we did with the two versions of Green's theorem and with Stokes' theorem, we can use the divergence theorem to deduce a geometric interpretation of the divergence of a vector field. You are encouraged to think about this.

7.1.2 Examples

Example 7.2. Let's use the divergence theorem to compute the flux of the vector field

$$\mathbf{F}(x, y, z) = (x - z)\mathbf{i} + (y - x)\mathbf{j} + (z - y)\mathbf{k}$$

through the boundary $\partial\mathcal{W}$ of the solid cylinder

$$\mathcal{W} = \{(x, y, z) : x^2 + y^2 \le a^2, 0 \le z \le 1\},$$

oriented with outward-pointing normal vector. Notice that $\partial\mathcal{W}$ has three components: the discs on top and bottom, as well as the hollow cylinder defined by $x^2 + y^2 = a^2$. So evaluating $\oiint_{\partial\mathcal{W}} \mathbf{F} \cdot d\mathbf{S}$ without the

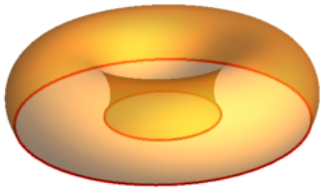

Figure 7.1: A surface $\mathcal{S}$ that we'd rather not parametrize.

divergence theorem would require that we compute three separate surface integrals. But according to the divergence theorem,

$$\oiint_{\partial \mathcal{W}} \mathbf{F} \cdot d\mathbf{S} = \iiint_{\mathcal{W}} \operatorname{div}(\mathbf{F})\, dV.$$

Now $\operatorname{div}(\mathbf{F}) = 1+1+1 = 3$, so the integral of $\operatorname{div}(\mathbf{F})$ over $\mathcal{W}$ is 3 $\operatorname{vol}(\mathcal{W}) = 3\pi a^2$. We conclude that

$$\boxed{\oiint_{\partial \mathcal{W}} \mathbf{F} \cdot d\mathbf{S} = 3\pi a^2.}$$

Example 7.3. Let's consider the surface $\mathcal{S}$ defined by

$$\mathcal{S} := \{(x,y,z) : (\sqrt{x^2+y^2}-4)^2 + (z-1)^2 = 4, z \geq 0\}.$$

This terrifying surface is seen in Figure 7.1. We want to compute the flux of $\mathbf{F}(x,y,z) := \langle 2xy, -y^2, 1\rangle$ through $\mathcal{S}$, with outward-pointing normal vector. We'd really like to avoid doing this directly, so our mind turns to the divergence theorem. Applying this should make our lives a good bit easier, since

$$\operatorname{div}(\mathbf{F}) = \frac{\partial}{\partial x}(2xy) + \frac{\partial}{\partial y}(-y^2) + \frac{\partial}{\partial z}(1) = 2y - 2y + 0 = 0.$$

The main problem we have is that $\mathcal{S}$ is not the boundary of any region in $\mathbb{R}^3$. But consider the region $\mathcal{W}$ defined by

$$\mathcal{W} := \{(x,y,z) : (\sqrt{x^2+y^2}-4)^2 + (z-1)^2 \leq 4, z \geq 0\}.$$

One component of $\partial\mathcal{W}$ is $\mathcal{S}$, and we let $\mathcal{A}$ denote the other component. This is the annulus seen at the bottom of $\mathcal{S}$ in Figure 7.1, and corresponds to setting $z = 0$. A little algebraic manipulation shows that

$$\mathcal{A} = \{(x,y,0) : (4-\sqrt{3})^2 \leq x^2+y^2 \leq (4+\sqrt{3})^2\},$$

so $\mathcal{A}$ is bounded by the circles of radius $4+\sqrt{3}$ and $4-\sqrt{3}$ centered at the origin in the xy-plane. Now the divergence theorem tells us that

$$0 = \iiint_{\mathcal{W}} \operatorname{div}(\mathbf{F})\, dV = \iint_{\mathcal{S}} \mathbf{F} \cdot d\mathbf{S} + \iint_{\mathcal{A}} \mathbf{F} \cdot d\mathbf{S},$$

where $\mathcal{S}$ has the outward-pointing normal vector and $\mathcal{A}$ has the downward-pointing normal vector (convince yourself that this is the correct orientation). So the oriented unit normal vector for $\mathcal{A}$ is $-\mathbf{k}$, and we see that $\mathbf{F} \cdot \mathbf{N} \equiv -1$ on $\mathcal{A}$. Altogether,

$$\iint_{\mathcal{S}} \mathbf{F} \cdot d\mathbf{S} = -\iint_{\mathcal{A}} \mathbf{F} \cdot d\mathbf{S} = -\iint_{\mathcal{A}} -1\, dS = \operatorname{area}(\mathcal{A}).$$

Now the area of an annulus is given by $\pi(R^2 - r^2)$, where R is the radius of the outer circle and r is the radius of the inner circle. So

$$\operatorname{area}(\mathcal{A}) = \pi\left[(4+\sqrt{3})^2 - (4-\sqrt{3})^2\right] = 16\pi\sqrt{3}.$$

Finally, we conclude that $\iint_{\mathcal{S}} \mathbf{F} \cdot d\mathbf{S} = \boxed{16\pi\sqrt{3}}$. This was a lot of work, but I think it's still better than a direct computation. In the end, we didn't have to find any antiderivatives.

7.2 Day 19: Connections to topology

Goals

By the end of today's class, we should be able to do the following.

1. Describe which open regions in $\mathbb{R}^2$ admit curl-free-but-not-conservative vector fields.
2. Discuss the relationship between the integral of $\mathbf{F}$ over $\mathcal{C}$ and the **topological type** of $\mathcal{C}$, where $\mathbf{F}$ is a **curl-free vector field** and $\mathcal{C}$ is a **closed loop** in $\mathbb{R}^2$.

A quick recap of the tour through integral vector calculus we've taken over the last few weeks is to say that we introduced some derivatives for vector fields, integrated them, and then stated some FTICs. Of course this description lacks nuance, but it hits the highlights. Today we want to back up and think about the derivatives grad, curl, and div again. Remember that we defined curl as a means of encoding a necessary condition for a vector field to be conservative: in $\mathbb{R}^3$ we have

$$\mathbf{F} = \operatorname{grad} \phi \quad \Rightarrow \quad \operatorname{curl}(\mathbf{F}) \equiv \mathbf{0}.$$

Similarly, divergence gives us a necessary condition for a vector field to admit a vector potential:

$$\mathbf{F} = \operatorname{curl} \mathbf{G} \quad \Rightarrow \quad \operatorname{div}(\mathbf{F}) \equiv 0.$$

Poincaré's lemma told us that if the domain of our vector field is contractible — for instance, if the domain is all of $\mathbb{R}^3$ — then these necessary conditions are actually sufficient. But we've also seen curl-free vector fields which don't admit scalar potentials and divergence-free vector fields which don't admit vector potentials. These vector fields necessarily have domains which are not contractible, and our goal today is to investigate what features of the domain allow these strange vector fields to exist.

7.2.1 How will we distinguish regions and curves?

Throughout today's discussion, we're going to refer to the **topology** of the domain of a vector field, or of a curve or surface over which we want to integrate the vector field. We won't be precise about what this means, but will instead consider two open regions $\mathcal{R}_1, \mathcal{R}_2 \subset \mathbb{R}^n$ to be *topologically equivalent* if we can continuously deform $\mathcal{R}_1$ into $\mathcal{R}_2$. For instance, we'll consider the punctured plane $\mathbb{R}^2 \setminus \{(0,0)\}$ to be topologically equivalent to the region we get by deleting a disc of some positive radius from $\mathbb{R}^2$. We continuously deform the former into the latter by gradually changing the size of the hole in our region.

Once we've fixed a region $\mathcal{R}$, we'll consider two parametrized curves in $\mathcal{R}$ to be topologically equivalent if we can continuously deform one parametrized curve into the other *within* $\mathcal{R}$. Some comments about this are

(1) the orientations of the curves matter, so a parametrized curve need not be topologically equivalent to the same curve with the opposite orientation;

(2) the continuous deformation must happen within $\mathcal{R}$, so a pair of curves can be topologically distinct in one region, but topologically equivalent when considered in a larger region;

(3) this notion of topological equivalence is reasonably close to **regular homotopy** (but we don't necessarily require our curves to be regular).

Even though an open region in $\mathbb{R}^n$ might be "continuously deformed" to a curve, we're not going to consider regions and curves to be topologically equivalent to one another: regions can only be equivalent to other regions, and curves to other curves.

We can write down a similar notion of topological equivalence for parametrized surfaces, but today we're going to focus on curves and regions in $\mathbb{R}^2$ — if only for the sake of drawing pictures.

Disclaimer. As a general remark, we're just telling a nice story today, and not being super careful about details. The details can get somewhat technical, so we leave this to a course on topology. The material we discuss today can be accepted as a story, but maybe not as precise mathematical facts. (Indeed, a lot of things won't even be stated precisely.)

7.2.2 Poincaré's lemma and integrals along closed loops

Back in Week 4, we stated Poincaré's lemma for vector fields defined on contractible subsets of $\mathbb{R}^3$. Let's write down Poincaré's lemma for vector fields on $\mathbb{R}^2$.

Theorem 7.4: Poincaré's lemma in $\mathbb{R}^2$

Suppose $A \subset \mathbb{R}^2$ is contractible and $\mathbf{F}: A \to \mathbb{R}^2$ is a C^1-vector field. Then $\mathbf{F}$ is conservative if and only if $\operatorname{curl}_z(\mathbf{F}) \equiv 0$ on A.

From the definition of curl_z, we knew one direction of this statement: if $\mathbf{F}$ is conservative, then $\operatorname{curl}_z(\mathbf{F}) = 0$. Poincaré's lemma gives us the converse: a curl-free vector field on a contractible region in fact must be conservative. We saw by way of example that this need not be true for regions which are not contractible.

Example 7.5. Consider the vector field $\mathbf{F}: \mathcal{R} \to \mathbb{R}^2$ defined by

$$\mathbf{F}(x, y) := \frac{1}{x^2 + y^2}\langle -y, x\rangle,$$

where $\mathcal{R} = \mathbb{R}^2 \setminus \{(0,0)\}$. We've previously computed that $\operatorname{curl}_z(\mathbf{F}) \equiv 0$, but we've also shown that $\mathbf{F}$ cannot be conservative. We refer to this vector field as a **vortex vector field**.

You may recall that we proved that the above vector field fails to be conservative by showing that it admits closed orbits: parametrized curves which are everywhere tangent to $\mathbf{F}$. The problem is actually somewhat simpler than that. Consider integrating a conservative vector field $\nabla\phi$ over a closed loop $\mathcal{C}$ in $\mathbb{R}^2$. If $\mathcal{C}$ is parametrized by $\mathbf{r}: [a, b] \to \mathbb{R}^n$, then

$$\oint_{\mathcal{C}} \nabla\phi \cdot d\mathbf{r} = \int_a^b \nabla\phi(\mathbf{r}(t)) \cdot \mathbf{r}'(t)\, dt = \int_a^b \frac{d}{dt}(\phi(\mathbf{r}(t)))\, dt = \phi(\mathbf{r}(b)) - \phi(\mathbf{r}(a)).$$

The second equality uses the definition of the gradient given in Week 4. But since $\mathbf{r}$ parametrizes the closed curve $\mathcal{C}$, we must have $\mathbf{r}(b) = \mathbf{r}(a)$. So the integral of $\nabla\phi$ over $\mathcal{C}$ is 0, and we have the following important fact:

Fact. The integral of a conservative vector field over a closed loop is 0.

We can verify that our vortex vector field $\mathbf{F}$ is not conservative with the following computation.

Exercise 7.6. Let $\mathcal{C}$ be the unit circle in $\mathbb{R}^2$ centered at $(0,0)$, with counterclockwise orientation, and let $\mathbf{F}$ be the vector field defined in Example 7.5. Show that $\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = 2\pi$, and conclude that $\mathbf{F}$ is not conservative.

In fact, one can show that conservative vector fields are precisely those which integrate to zero over any[1] closed curve.

Exercise 7.7. Show that a vector field $\mathbf{F}: \mathcal{R} \to \mathbb{R}^n$, where $\mathcal{R}$ is a region in $\mathbb{R}^n$, is conservative if and only if $\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = 0$ for any closed curve $\mathcal{C} \subset \mathcal{R}$.

We point out that our curl-free-but-not-conservative vector field is only allowed to exist because our domain $\mathbb{R}^2 \setminus \{(0,0)\}$ has a hole. In a region $\mathcal{R}$ which has no holes (such as $\mathbb{R}^2$), every simple closed curve $\mathcal{C}$ is going to enclose a subregion $\mathcal{R}'$. If we have a curl-free vector field $\mathbf{F}$, then we can use Green's theorem to deduce that the integral of $\mathbf{F}$ over $\mathcal{C}$ is zero:

$$\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = \iint_{\mathcal{R}'} \operatorname{curl}_z(\mathbf{F})\, dA = 0.$$

If we want to produce a curl-free-but-not-conservative vector field on $\mathcal{R}$, we need $\mathbf{F}$ to integrate to a nonzero number over some simple closed curve $\mathcal{C}$ in $\mathcal{R}$, and now we see that this can only happen if $\mathcal{C}$ fails to bound a subregion of $\mathcal{R}$. In $\mathbb{R}^2$, having closed loops which fail to bound subregions corresponds precisely with a failure of **simple connectivity**.

Definition. We'll say that an open region $\mathcal{R}$ in $\mathbb{R}^2$ is **simply connected** if every closed loop in $\mathcal{C}$ can be

[1] Maybe we need to use a word like "rectifiable" here to make things make sense, but see the disclaimer.

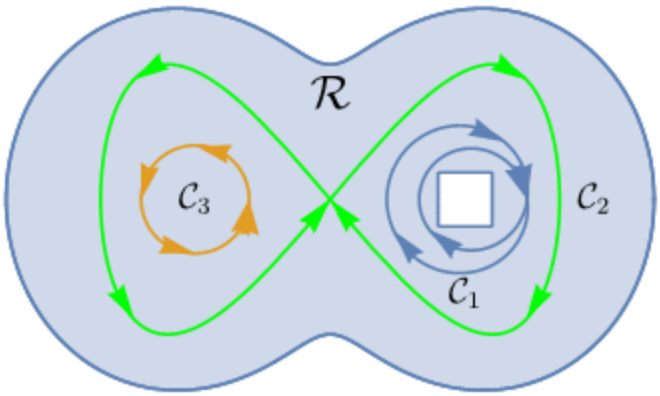

Figure 7.2: Curves $\mathcal{C}_1$, $\mathcal{C}_2$, and $\mathcal{C}_3$ have winding numbers -2, -1, and 0, respectively, about the square hole in $\mathcal{R}$.

continuously deformed to a point.

Said another way, a region in $\mathbb{R}^2$ is simply connected if it is connected and contractible. So Poincaré's lemma holds for simply connected regions, but for non-simply connected regions we can produce curl-free vector fields which evaluate to nonzero numbers on closed loops. Loosely, we do this by writing down a vortex vector field for each hole in the region.

7.2.3 Curl-free vector fields care only about topological type

Okay, so our ability to write down curl-free vector fields which are not conservative corresponds to the existence of holes in our region $\mathcal{R}$. Let's say we have an open region $\mathcal{R}$ in $\mathbb{R}^2$ which is connected and has n holes, for some $n \geq 0$. A natural question then asks how many different curl-free-but-not-conservative vector fields $\mathcal{R}$ admits. Let's delay this question for a moment and start with something slightly simpler.

Our first goal is to show that if $\mathbf{F}$ is a curl-free vector field on $\mathcal{R}$ and $\mathcal{C}$ is an oriented closed curve in $\mathcal{R}$, then $\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r}$ depends only on the topological type of $\mathcal{C}$. That is, if we make continuous deformations to $\mathcal{C}$, we won't change the value of this integral. Our argument will necessarily be very informal.

Before we can state our claim, we need to describe the notion of **winding number**. As with so many things we discuss today, this is something that could be defined formally, but for which we'll just give an informal idea.

Definition. Let p be a point in $\mathbb{R}^2$, and let $\mathcal{C}$ be a closed, oriented (but not necessarily simple) curve in $\mathbb{R}^2$ which does not pass through p. We obtain the **winding number of $\mathcal{C}$ about p** by counting the number of times $\mathcal{C}$ winds about p in a counterclockwise manner. This is a **signed count**, so that clockwise rotations of $\mathcal{C}$ about p count negatively.

We will speak informally of the *holes* in our region $\mathcal{R}$; the winding number of a closed, oriented curve $\mathcal{C}$ about a hole is simply the winding number of $\mathcal{C}$ about any point in the hole.

Example 7.8. Consider the region $\mathcal{R}$ in Figure 7.2. This region has one hole, and we can easily compute the winding numbers of a curve about this hole. Because $\mathcal{C}_1$ winds twice clockwise about the hole, its winding number is -2; the curve $\mathcal{C}_2$ has some interesting behavior away from the hole, but ultimately makes one counterclockwise rotation about the hole, so its winding number is 1; finally, $\mathcal{C}_3$ does not wind around the hole at all, so its winding number is 0.

Certainly the winding number isn't affected by continuous deformations; our plan is to show that if $\text{curl}_z(\mathbf{F}) \equiv 0$, then $\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r}$ depends only on the winding numbers of $\mathcal{C}$ about the various holes of $\mathcal{R}$. To this end, let's label the holes of $\mathcal{R}$ with the names $H_1, \ldots, H_n$ and choose pairwise disjoint, simple, closed, oriented curves $\mathcal{C}_1, \ldots, \mathcal{C}_n$ so that the winding number of $\mathcal{C}_i$ about H_j is 1 if $i = j$ and 0 if $i \neq j$. Morally, we want $\mathcal{C}_k$ to

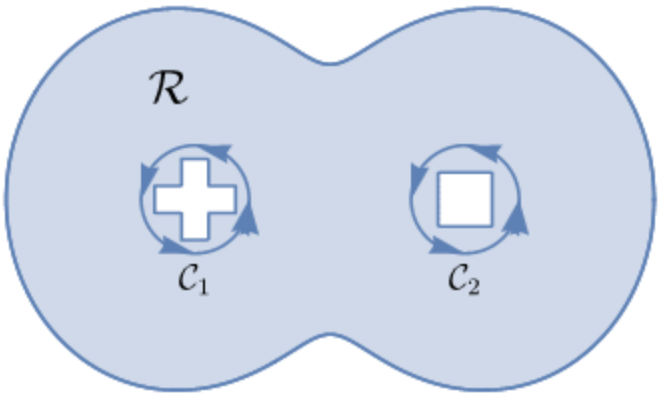

Figure 7.3: The curves $\mathcal{C}_1$ and $\mathcal{C}_2$ are "boundary parallel."

be the boundary of the hole H_k, oriented counterclockwise. The problem is that ∂H_k need not be very nice — for instance, if our hole is just a puncture, then ∂H_k is empty. Figure 7.3 gives an example where $n = 2$.

Now let's fix a C^1-vector field $\mathbf{F}\colon \mathcal{R} \to \mathbb{R}^2$ which has $\text{curl}_z(\mathbf{F}) \equiv 0$ on $\mathcal{R}$. For $1 \leq k \leq n$ we define

$$I_k := \oint_{\mathcal{C}_k} \mathbf{F} \cdot d\mathbf{r}.$$

The moral of our earlier discussion is that because $\mathcal{C}_k$ encloses a hole in $\mathcal{R}$ (rather than a subregion of $\mathcal{R}$), this integral need not be zero. At last we make the following (still-not-quite-precise) claim about the integral of $\mathbf{F}$ over an arbitrary closed curve $\mathcal{C}$ in $\mathcal{R}$.

Proposition 7.9. *Let $\mathcal{C} \subset \mathcal{R}$ be an oriented, closed curve, and let $\mathbf{F}\colon \mathcal{R} \to \mathbb{R}$ be a curl-free vector field on an open region $\mathcal{R} \subset \mathbb{R}^2$. Suppose that $\mathcal{R}$ has n holes, with curves $\mathcal{C}_1, \ldots, \mathcal{C}_n$ as identified above, and define the values $I_1, \ldots, I_n \in \mathbb{R}$ as above. Then*

$$\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = w_1 I_1 + \cdots + w_n I_n,$$

where w_k is the winding number of $\mathcal{C}$ about the k^{th} hole of $\mathcal{R}$.

Since we haven't bothered to precisely define all of the words in our proposition, we certainly won't attempt a precise proof, but we can give an informal argument. In fact, we won't even try to prove the claim for an arbitrary curve $\mathcal{C}$, but instead just consider some cases and claim that the main ideas of the general argument are all covered by these cases.

Case 1. First, let's suppose that $\mathcal{C}$ is simple, oriented counterclockwise, and encloses the holes $H_{k_1}, \ldots, H_{k_m}$, for some collection of integers $\{k_1, \ldots, k_m\} \subset \{1, \ldots, n\}$. Figure 7.4 shows an example where $\mathcal{C}$ encloses holes H_1 and H_2. We also assume that $\mathcal{C}$ does not intersect any of the curves $\mathcal{C}_{k_1}, \ldots, \mathcal{C}_{k_m}$, though such intersections don't create much trouble. What we should see from Figure 7.4 is that there is a subregion $\mathcal{R}'$ of $\mathcal{R}$ which is bounded by $\mathcal{C}$ and the curves $\mathcal{C}_{k_1}, \ldots, \mathcal{C}_{k_m}$. Of course the curves $\mathcal{C}_{k_1}, \ldots, \mathcal{C}_{k_m}$ don't carry the boundary orientation induced by $\mathcal{R}'$, so we can think of this boundary as

$$\partial \mathcal{R}' = \mathcal{C} \sqcup (-\mathcal{C}_{k_1}) \sqcup \cdots \sqcup (-\mathcal{C}_{k_m}).$$

We know that $\mathbf{F}$ is curl-free, so we can apply Green's theorem to $\mathbf{F}$ over the region $\mathcal{R}'$ to find that

$$\begin{aligned} 0 &= \iint_{\mathcal{R}'} \text{curl}_z(\mathbf{F})\, dA = \oint_{\partial \mathcal{R}'} \mathbf{F} \cdot d\mathbf{r} = \oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} - \oint_{\mathcal{C}_{k_1}} \mathbf{F} \cdot d\mathbf{r} - \cdots - \oint_{\mathcal{C}_{k_m}} \mathbf{F} \cdot d\mathbf{r} \\ &= \oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} - I_{k_1} - \cdots - I_{k_m}. \end{aligned}$$

Rearranging this leads us to the conclusion that

$$\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = I_{k_1} + \cdots + I_{k_m},$$

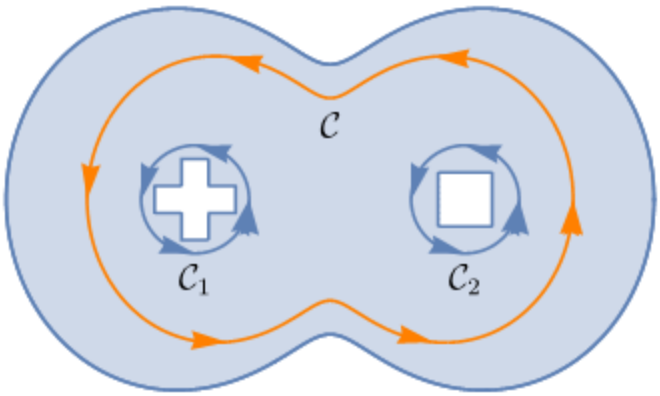

Figure 7.4: The curve $\mathcal{C}$ has winding number 1 about each hole.

which matches the conclusion of Proposition 7.9, since the winding number of $\mathcal{C}$ about H_{k_i} is 1, while the winding number of $\mathcal{C}$ about holes not enclosed by $\mathcal{C}$ is 0.

Case 2. Next, let's consider the case where $\mathcal{C}$ is no longer required to be simple, but encloses just one hole H_k. This case is depicted in Figure 7.5a. Notice that we no longer require $\mathcal{C}$ to carry a counterclockwise orientation, but we once again make the mild assumption that $\mathcal{C}$ does not intersect $\mathcal{C}_k$. The trick this time is to break $\mathcal{C}$ into several simple closed curves, as seen in Figure 7.5b. We can classify these subcurves into (a) those which enclose H_k with a counterclockwise orientation; (b) those which enclose H_k with a clockwise orientation; (c) and those which don't enclose H_k at all. It should be reasonably believable that the winding number of $\mathcal{C}$ about H_k is then given by the number of subcurves of type (a) minus the number of type (b). Green's theorem will pretty easily show that the integral of $\mathbf{F}$ over subcurves of type (c) is 0. We can then analyze the subcurves of types (a) and (b) as in Case 1. What we'll find is that each subcurve of type (a) contributes I_k to our integral, while each subcurve of type (b) contributes $-I_k$. The upshot is that

$$\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = (b-a)I_k,$$

where a and b are number of subcurves of types (a) and (b), respectively. But $w_k = b-a$, so we have the conclusion of Proposition 7.9!

Case 3. Finally, we consider a curve $\mathcal{C}$ which is allowed to enclose holes $H_{k_1}, \ldots, H_{k_m}$ and is not required to be simple. For instance, see Figure 7.6. We compute $\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r}$ in much the same way we did in Case 2. We start by breaking $\mathcal{C}$ into simple closed curves $\mathcal{D}_1, \ldots, \mathcal{D}_N$. For each of these, we compute $\oint_{\mathcal{D}_i} \mathbf{F} \cdot d\mathbf{r}$ as we did in

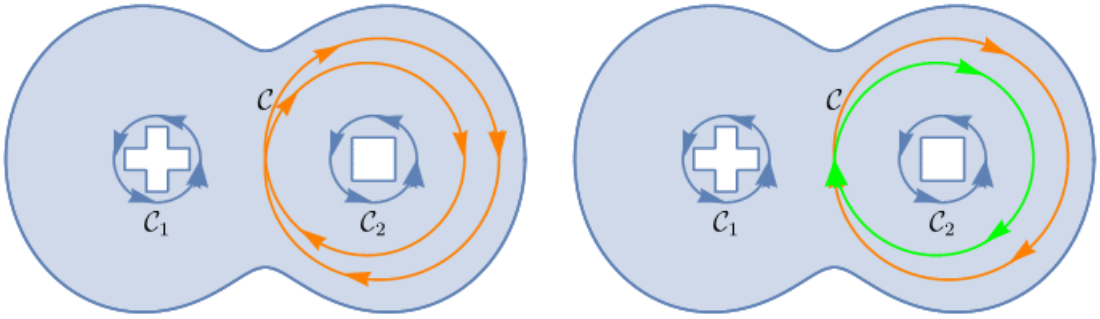

(a) A curve with winding number -2.

(b) A pair of curves, each with winding number -1.

Figure 7.5: A closed curve which fails to be simple can be broken into simple closed curves.

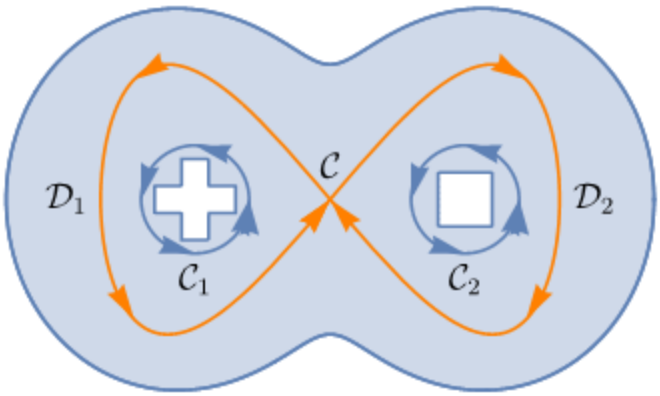

Figure 7.6: The curve $\mathcal{C}$ is not simple and encloses multiple holes.

Case 1: we consider a region bounded on one side by $\mathcal{D}_i$, and on the other by the curves which correspond to the hole(s) enclosed by $\mathcal{D}_i$. After accounting for orientations, we should find that

$$\oint_{\mathcal{D}_i} \mathbf{F} \cdot d\mathbf{r} = w_1(\mathcal{D}_i)I_1 + \cdots + w_n(\mathcal{D}_i)I_n, \tag{7.3}$$

where $w_j(\mathcal{D}_i)$ is the winding number of $\mathcal{D}_i$ about the hole H_j. It's probably not hard to believe that the winding numbers of $\mathcal{C}$ can be obtained by adding up those of the constituent subcurves:

$$w_j(\mathcal{C}) = w_j(\mathcal{D}_1) + \cdots + w_j(\mathcal{D}_N).$$

Then we get the desired conclusion by summing the equations (7.3) as i varies from 1 to N.

Let's recap the result we've sketched here. We want to show that if $\mathbf{F}\colon \mathcal{R} \to \mathbb{R}^2$ is a curl-free vector field and $\mathcal{C}$ is a closed curve in $\mathcal{R}$, then $\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r}$ depends only on the topological type of $\mathcal{C}$. That is, the value given by this integral does not change when we make continuous deformations to $\mathcal{C}$. Our argument in support of this claim showed that $\oint_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r}$ could be computed in terms of the winding numbers of $\mathcal{C}$ about the various holes of $\mathcal{R}$. We then blackboxed the fact that the winding numbers are unchanged by continuous deformation, since this fact seems clear enough from pictures. In summary: curl-free vector fields care only about topological type.

7.2.4 Other topological considerations

The above discussion, while interesting, leaves mostly unexplored the question of just how many curl-free-but-not-conservative vector fields we can write down on a connected, open region $\mathcal{R}$ in $\mathbb{R}^2$. Of course the answer is either 0 (if $\mathcal{R}$ is simply connected) or infinitely many (if $\mathcal{R}$ has at least one hole): once we have one vector field which works, we can scale it by any nonzero real number. What we're really asking is the somewhat more nebulous question of how many *independent* curl-free-but-not-conservative vector fields $\mathcal{R}$ will admit. We won't address this in detail here, but the loose answer is that we get one such vector field $\mathbf{F}_k$ for each hole H_k in $\mathcal{R}$, and any curl-free vector field other than these can be written as a linear combination of $\mathbf{F}_1, \ldots, \mathbf{F}_n$, plus the conservative vector fields on $\mathcal{R}$. This answer can be made more precise once we have a bit more linear algebra under our belts.

There are also a couple of generalizations of the story we've told today that are interesting. For instance, for an open region $\mathcal{R}$ in $\mathbb{R}^3$, we can ask for curl-free-but-not-conservative vector fields, and we can also ask for vector fields which are divergence-free, but do not admit a vector potential. As in two dimensions, the existence of such vector fields is dictated by the topology of $\mathcal{R}$. We again find curl-free-but-not-conservative vector fields when $\mathcal{R}$ fails to be simply connected, and divergence-free-but-without-vector-potential vector fields arise when some two-dimensional version of simple connectivity fails. Basically, the former vector fields exist when we have deleted a line, and the latter exist when we delete a hole. One can then show that integration of these vector fields over closed curves or surfaces, respectively, depends only on the topological

type of the underlying curve or surface.

We can also ask these questions for vector fields on surfaces. Say we have a surface $\mathcal{S}$ and a vector field $\mathbf{F}$ for which curl($\mathbf{F}$) is always tangent to $\mathcal{S}$. Does this mean that we can find a function $\phi : \mathbb{R}^3 \to \mathbb{R}$ such that $\text{proj}_{T_p\mathcal{S}}\mathbf{F} \equiv \nabla(\phi|_{\mathcal{S}})$? If not, can we at least say that the integral of $\mathbf{F}$ over a closed curve $\mathcal{C}$ depends only on the topological type of $\mathcal{C}$? Given a surface $\mathcal{S} \subset \mathbb{R}^3$, how many vector fields can we write down whose curl is everywhere tangent to $\mathcal{S}$, but which are not conservative? You are invited and encouraged to think about all of these questions.

Exercise 7.10. Think about — and maybe even answer — the various questions raised in this subsection.

8 Week 8

We'll kick our week off with a midterm and follow that up with a step towards the abstract. The midterm wraps up part two of our course, and we then begin part three, which will attempt to generalize some of the key points of part two to any dimension. Ultimately, the goal of this last portion of the course is to prove a **generalized Stokes' theorem**. We've described Green's theorems, Stokes' theorem (for surface integrals), and the divergence theorems as different versions of the Fundamental Theorem of Integral Calculus. What we want now is a version of these results which holds in all dimensions, and from which we may recover these concrete theorems — one FTIC to rule them all!

There is a lot of road we must cover before we can even state the generalized Stokes' theorem that we're after. FTICs come in the form of an integral equation, and allow us to trade boundaries for derivatives, so we'll need to identify:

(1) the objects which will serve as our integrands;

(2) the domains over which we will integrate;

(3) how to differentiate our integrands;

(4) how to determine the boundaries of our domains.

In our two class meetings this week (after the midterm), we'll address item (1) by introducing **differential forms** and some of the algebraic operations that can be performed on them.

As a final note before we begin, let's save ourselves some adjectives by making the following assumption.

Everything is smooth. For the duration of this course, we don't want to worry about just how differentiable a map is — such considerations are somewhat beside the points we want to make. Thus, unless otherwise stated, all maps of any sort (scalar functions, vector fields, coordinate systems, et cetera) will be assumed to be *infinitely differentiable*, or *smooth*.

8.1 Day 20: Midterm 2

Goals

By the end of today's class, we should be able to do the following.

1. Solve the problems on the midterm.

Here are the midterm 2 problems and solutions.

1. The purpose of this problem is to prove that conservative vector fields are precisely those vector fields[1] which are *path-independent*. That is, we will show that $\mathbf{F}\colon A \to \mathbb{R}^n$ is conservative if and only if, for any two points $P, Q \in A$, the line integral of $\mathbf{F}$ over an oriented curve from P to Q is independent of the choice of curve. Here A is an open subset of $\mathbb{R}^n$.

   (a) Suppose that $\mathbf{F}$ is conservative, and write $\mathbf{F} = \nabla\phi$ for some sufficiently differentiable function $\phi : A \to \mathbb{R}$. Let $\mathcal{C} \subset A$ be any oriented, regular curve from P to Q. Prove that

$$\int_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = \phi(Q) - \phi(P).$$

Solution. We begin by choosing a simple, regular parametrization $\mathbf{r}\colon [a,b] \to A$ of $\mathcal{C}$, with $\mathbf{r}(a) = P$ and $\mathbf{r}(b) = Q$. Then we have

$$\int_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = \int_a^b \mathbf{F}(\mathbf{r}(t)) \cdot \mathbf{r}'(t)\, dt = \int_a^b \nabla\phi(\mathbf{r}(t)) \cdot \mathbf{r}'(t)\, dt.$$

[1] Let's say that *vector field* means *infinitely differentiable vector field*, so that we don't have to worry about C^1 versus C^2 versus whatever.

According to the definition of the gradient, $\nabla\phi(\mathbf{r}(t))\cdot\mathbf{r}'(t) = \frac{d}{dt}(\phi(\mathbf{r}(t)))$. So we have

$$\begin{aligned}\int_{\mathcal{C}} \mathbf{F}\cdot d\mathbf{r} &= \int_a^b \frac{d}{dt}(\phi(\mathbf{r}(t)))\,dt = [\phi(\mathbf{r}(t))]_a^b \\ &= \phi(\mathbf{r}(b)) - \phi(\mathbf{r}(a)) = \phi(Q) - \phi(P),\end{aligned}$$

as desired. So $\mathbf{F}$ is path-independent. □

(b) Conversely, suppose that $\mathbf{F}$ is path-independent. That is, assume that if $\mathcal{C}_1, \mathcal{C}_2 \subset A$ are oriented curves with the same oriented boundary, then $\int_{\mathcal{C}_1}\mathbf{F}\cdot d\mathbf{r} = \int_{\mathcal{C}_2}\mathbf{F}\cdot d\mathbf{r}$. Now define $\phi: A\to\mathbb{R}$ as follows. Fix an arbitrary point $P\in A$ and declare $\phi(P) = 0$. For any $Q\in A$, define

$$\phi(Q) := \int_{\mathcal{C}} \mathbf{F}\cdot d\mathbf{r},$$

where $\mathcal{C}$ is an oriented curve in A from P to Q. (We are assuming that A is *path-connected*, meaning that any two points in A may be connected by a (regular) curve.)

i. Prove that ϕ is a well-defined function. (That is, each input $Q\in A$ has only one output.)

Solution. What we need to show here is that the choice of curve $\mathcal{C}$ made in the definition of $\phi(Q)$ does not affect the outcome. But this is precisely what path-independence guarantees. Namely, if $\mathcal{C}'$ is some other curve in A from P to Q, then we know that

$$\int_{\mathcal{C}'}\mathbf{F}\cdot d\mathbf{r} = \int_{\mathcal{C}}\mathbf{F}\cdot d\mathbf{r},$$

and thus we may define $\phi(Q)$ using either curve. □

ii. Use the Fundamental Theorem of Calculus to prove that $\mathbf{F} = \nabla\phi$.

Solution. We think about the derivative of $\phi(\mathbf{r}(t))$, where $\mathbf{r}: I\to A$ is an arbitrary regular parametrized curve in A. We can extend our parametrized curve so that $\mathbf{r}(t_0) = P$ for some particular value t_0. This means that $\mathbf{r}$ parametrizes a curve from P to $\mathbf{r}(t)$, for any $t\in I$, and thus we have

$$\phi(\mathbf{r}(t)) = \int_{t_0}^{t}\mathbf{F}(\mathbf{r}(u))\cdot\mathbf{r}'(u)\,du.$$

(We use the variable u since t is already in use.) We compute the derivative using the Fundamental Theorem of Calculus:

$$\frac{d}{dt}(\phi(\mathbf{r}(t))) = \frac{d}{dt}\left(\int_{t_0}^{t}\mathbf{F}(\mathbf{r}(u))\cdot\mathbf{r}'(u)\,du\right) = \mathbf{F}(\mathbf{r}(t))\cdot\mathbf{r}'(t).$$

But this equation which is satisfied by $\mathbf{F}$ is precisely the equation we expect the gradient of ϕ to satisfy! So we conclude that $\mathbf{F} = \nabla\phi$. □

2. The **open upper half plane** $\mathbb{H}$ is the subset of $\mathbb{R}^2$ defined by

$$\mathbb{H} := \{(x,y) : y > 0\}.$$

Given a region $\mathcal{R}\subset\mathbb{H}$, the **hyperbolic area** of $\mathcal{R}$ is defined to be:

$$\text{hArea}(\mathcal{R}) := \iint_{\mathcal{R}}\frac{1}{y^2}\,dA,$$

where dA is the usual Euclidean area form.

(a) Use Green's theorem to rewrite the hyperbolic area of $\mathcal{R}$ as

$$\text{hArea}(\mathcal{R}) = \oint_{\partial\mathcal{R}}\mathbf{F}\cdot d\mathbf{r}$$

for some vector field $\mathbf{F}$.

Solution. We want a vector field $\mathbf{F} = \langle F_1, F_2 \rangle$ which has $\operatorname{curl}_z(\mathbf{F}) = y^{-2}$. Given such a vector field, Green's theorem tells us that

$$\oint_{\partial \mathcal{R}} \mathbf{F} \cdot d\mathbf{r} = \iint_{\mathcal{R}} \operatorname{curl}_z(\mathbf{F})\, dA = \iint_{\partial \mathcal{R}} \frac{1}{y^2}\, dA = \operatorname{hArea}(\mathcal{R}).$$

So we're trying to solve the partial differential equation

$$\frac{1}{y^2} = \operatorname{curl}_z(\mathbf{F}) = \frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y}.$$

As suggested by the hint, we take $F_2 \equiv 0$, so that we're now trying to solve $\frac{\partial F_1}{\partial y} = -y^{-2}$. But this is easy enough: we can just take $F_1(x, y) = y^{-1}$. So we find that

$$\boxed{\operatorname{hArea}(\mathcal{R}) = \oint_{\partial \mathcal{R}} \left\langle y^{-1}, 0 \right\rangle \cdot d\mathbf{r}.}$$

□

(b) Consider the region $\mathcal{T} \subset \mathbb{H}$ depicted here:

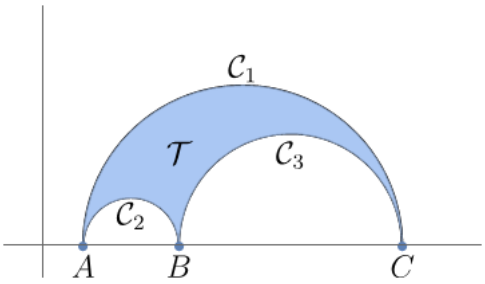

Each of the curves $\mathcal{C}_1$, $\mathcal{C}_2$, and $\mathcal{C}_3$ is a semicircle whose center is on the x-axis; $\mathcal{C}_1$ and $\mathcal{C}_2$ both pass through A; $\mathcal{C}_2$ and $\mathcal{C}_3$ both pass through B; and $\mathcal{C}_3$ and $\mathcal{C}_1$ both pass through C. Use your work from part (2a) to show that the hyperbolic area of $\mathcal{T}$ is π.

Solution. Let's start by proving the following claim:

Claim. If $\mathcal{C}$ is a semicircle in $\mathbb{H}$ centered on the x-axis and oriented counterclockwise, then $\int_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} = -\pi$, where $\mathbf{F} = \langle y^{-1}, 0 \rangle$.

Indeed, such a semicircle $\mathcal{C}$ admits the parametrization

$$\mathbf{r}(t) = (c + R\cos t, R\sin t), \quad 0 \le t \le \pi,$$

where R and $(c, 0)$ are the radius and center of $\mathcal{C}$, respectively. Having been motivated by polar coordinates, this parametrization is evidently simple; by observing that $\|\mathbf{r}'(t)\| \equiv R$, we see that $\mathbf{r}$ is regular. Then

$$\begin{aligned}
\int_{\mathcal{C}} \mathbf{F} \cdot d\mathbf{r} &= \int_0^\pi \mathbf{F}(\mathbf{r}(t)) \cdot \mathbf{r}'(t)\, dt \\
&= \int_0^\pi \langle (R\sin t)^{-1}, 0 \rangle \cdot \langle -R\sin t, R\cos t \rangle\, dt \\
&= \int_0^\pi -1\, dt = -\pi.
\end{aligned}$$

Given this claim (and our work in part (2a)), we can easily compute $\operatorname{hArea}(\mathcal{T})$:

$$\operatorname{hArea}(\mathcal{T}) = \oint_{\partial \mathcal{T}} \langle y^{-1}, 0 \rangle \cdot d\mathbf{r} = \int_{\mathcal{C}_1} \langle y^{-1}, 0 \rangle \cdot d\mathbf{r} + \int_{\mathcal{C}_2} \langle y^{-1}, 0 \rangle \cdot d\mathbf{r} + \int_{\mathcal{C}_3} \langle y^{-1}, 0 \rangle \cdot d\mathbf{r},$$

where $\mathcal{C}_1$ is oriented counterclockwise, while $\mathcal{C}_2$ and $\mathcal{C}_3$ are oriented clockwise. From this it follows that

$$\text{hArea}(\mathcal{T}) = -\pi + \pi + \pi = \boxed{\pi.}$$

□

3. Let $\mathbf{F}, \mathbf{G}\colon \mathbb{R}^3 \to \mathbb{R}^3$ be defined by

$$\mathbf{F}(x,y,z) = \langle y^2, 2z+x, 2y^2\rangle \quad \text{and} \quad \mathbf{G}(x,y,z) = \langle xz, \tfrac{1}{2}x^2, 0\rangle.$$

Find an equation for a plane P in $\mathbb{R}^3$ which passes through $(0,0,0)$ and has the property that

$$\oint_{\partial\mathcal{R}} \mathbf{F}\cdot d\mathbf{r} = 0 \quad \text{and} \quad \oint_{\partial\mathcal{R}} \mathbf{G}\cdot d\mathbf{r} = 0$$

for every region[2] $\mathcal{R}$ lying in P.

Solution. Since P passes through $(0,0,0)$, it can be described by an equation of the form $ax + by + cz = 0$, for some $a, b, c \in \mathbb{R}$ not all equal to 0. Moreover, the vector $\mathbf{n} = \langle a, b, c\rangle$ is then perpendicular to P. Now suppose that both curl($\mathbf{F}$) and curl($\mathbf{G}$) are perpendicular to $\mathbf{n}$, and let $\mathcal{R}$ be some (reasonably nice) region in P. Then, according to Stokes' theorem,

$$\oint_{\partial\mathcal{R}} \mathbf{F}\cdot d\mathbf{r} = \iint_{\mathcal{R}} \text{curl}(\mathbf{F})\cdot d\mathbf{S} = \iint_{\mathcal{R}} (\text{curl}(\mathbf{F})\cdot\mathbf{N})\, dS = 0,$$

and

$$\oint_{\partial\mathcal{R}} \mathbf{G}\cdot d\mathbf{r} = \iint_{\mathcal{R}} \text{curl}(\mathbf{G})\cdot d\mathbf{S} = \iint_{\mathcal{R}} (\text{curl}(\mathbf{G})\cdot\mathbf{N})\, dS = 0.$$

So we'll get the desired properties if we choose $\mathbf{n}$ (and thus $\mathbf{N}$) to be perpendicular to both curl($\mathbf{F}$) and curl($\mathbf{G}$), for all points (x,y,z). Let's compute curl($\mathbf{F}$) and curl($\mathbf{G}$). We have

$$\text{curl}(\mathbf{F}) = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \partial/\partial x & \partial/\partial y & \partial/\partial z \\ y^2 & 2z+x & 2y^2 \end{vmatrix} = \langle 4y-2, 0, 1-2y\rangle = (2y-1)\langle 2,0,-1\rangle$$

and

$$\text{curl}(\mathbf{G}) = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \partial/\partial x & \partial/\partial y & \partial/\partial z \\ xz & x^2/2 & 0 \end{vmatrix} = \langle 0, x, x\rangle = x\langle 0,1,1\rangle.$$

So the vectors curl($\mathbf{F}$) and curl($\mathbf{G}$) are not constant, but they are always parallel to $\langle 2,0,-1\rangle$ and $\langle 0,1,1\rangle$, respectively. Thus they are both perpendicular to

$$\langle 2,0,-1\rangle \times \langle 0,1,1\rangle = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ 2 & 0 & -1 \\ 0 & 1 & 1 \end{vmatrix} = \langle 1,-2,2\rangle$$

at all points (x,y,z). So we can take $\mathbf{n} = \langle 1,-2,2\rangle$, and thus we see that if P has equation $\boxed{x-2y+2z=0}$, then P passes through $(0,0,0)$ and has

$$\oint_{\partial\mathcal{R}} \mathbf{F}\cdot d\mathbf{r} = 0 \quad \text{and} \quad \oint_{\partial\mathcal{R}} \mathbf{G}\cdot d\mathbf{r} = 0$$

for every region $\mathcal{R}$ lying in P. □

4. In this problem we consider the vector field $\mathbf{F}\colon \mathbb{R}^3 \setminus \{(0,0,0)\} \to \mathbb{R}^3$ defined by

$$\mathbf{F}(x,y,z) = \frac{1}{(x^2+y^2+z^2)^{3/2}}\langle x,y,z\rangle.$$

[2] Assume that $\partial\mathcal{R}$ is a collection of nice curves, oriented using the boundary orientation. You can also assume that $\mathcal{R}$ admits a nice parametrization, but this isn't needed.

Our goal is to compute $\iint_{\partial\mathcal{W}} \mathbf{F}\cdot d\mathbf{S}$, where $\mathcal{W}$ is the region defined by

$$\mathcal{W} := \left\{(x,y,z) : (\sqrt{(x-R)^2+y^2}-R)^2+z^2 \le r^2\right\},$$

for some numbers $0 < r < R$. It's possible to compute this integral by parametrizing $\partial\mathcal{W}$, but we'll use the divergence theorem.

(a) Show that the origin is contained in $\mathcal{W}$, but not in $\partial\mathcal{W}$.
Note: Once we know this, we can find a number $\rho > 0$ so that the ball $\mathcal{B}_\rho$ with radius ρ centered at $(0,0,0)$ is contained in $\mathcal{W}$, but does not touch $\partial\mathcal{W}$.

Solution. Notice that $(0,0,0)$ satisfies the defining inequality of $\mathcal{W}$:

$$(\sqrt{(0-R)^2+0^2}-R)^2+0^2 = (\sqrt{R^2}-R)^2 = 0 \le r^2,$$

so $(0,0,0)\in\mathcal{W}$. But the boundary $\partial\mathcal{W}$ is defined by the equation

$$(\sqrt{(x-R)^2+y^2}-R)^2+z^2 = r^2,$$

which $(0,0,0)$ clearly does not satisfy. So $(0,0,0)$ is in $\mathcal{W}$ but not $\partial\mathcal{W}$. □

(b) Let $\mathcal{R}$ be the region obtained by deleting $\mathcal{B}_\rho$ from $\mathcal{W}$, where ρ is as chosen in part (4a). Show that $\operatorname{div}(\mathbf{F}) \equiv 0$ on $\mathcal{R}$.

Solution. Wherever $\mathbf{F}$ and its various derivatives are defined, the divergence of $\mathbf{F}$ is given by

$$\begin{aligned}
\operatorname{div}(\mathbf{F}) &= \frac{\partial}{\partial x}\left(\frac{x}{(x^2+y^2+z^2)^{3/2}}\right)+\frac{\partial}{\partial y}\left(\frac{x}{(x^2+y^2+z^2)^{3/2}}\right)+\frac{\partial}{\partial z}\left(\frac{x}{(x^2+y^2+z^2)^{3/2}}\right)\\
&= \frac{(x^2+y^2+z^2)^{3/2}-3x^2(x^2+y^2+z^2)^{1/2}}{(x^2+y^2+z^2)^3}+\frac{(x^2+y^2+z^2)^{3/2}-3y^2(x^2+y^2+z^2)^{1/2}}{(x^2+y^2+z^2)^3}\\
&\quad+\frac{(x^2+y^2+z^2)^{3/2}-3z^2(x^2+y^2+z^2)^{1/2}}{(x^2+y^2+z^2)^3}\\
&= \frac{3(x^2+y^2+z^2)^{3/2}-3(x^2+y^2+z^2)(x^2+y^2+z^2)^{1/2}}{(x^2+y^2+z^2)^3} = 0.
\end{aligned}$$

Since $\mathbf{F}$ is defined everywhere except $(0,0,0)$, the same is true of $\operatorname{div}(\mathbf{F})$. Finally, since the origin is in $\mathcal{B}_\rho$, this point is not in $\mathcal{R}$, and thus $\operatorname{div}(\mathbf{F})$ is defined and equal to 0 everywhere within $\mathcal{R}$. □

(c) Use the divergence theorem to prove that

$$\oiint_{\partial\mathcal{W}} \mathbf{F}\cdot d\mathbf{S} = \oiint_{\partial\mathcal{B}_\rho} \mathbf{F}\cdot d\mathbf{S}.$$

Each surface uses the corresponding boundary orientation.

Solution. The boundary of $\mathcal{R}$ has two components: $\partial\mathcal{W}$ and $\partial\mathcal{B}_\rho$. The boundary orientation of $\mathcal{R}$ gives $\partial\mathcal{W}$ the outward-pointing normal vector (as would $\mathcal{W}$), but the orientation on $\partial\mathcal{B}_\rho$ is inward-pointing, because the normal vector which points into $\mathcal{B}_\rho$ points out of $\mathcal{R}$. So the divergence theorem tells us that

$$\iiint_{\mathcal{R}} \operatorname{div}(\mathbf{F})\, dV = \oiint_{\partial\mathcal{W}} \mathbf{F}\cdot d\mathbf{S} - \oiint_{\partial\mathcal{B}_\rho} \mathbf{F}\cdot d\mathbf{S},$$

where the surfaces over which we integrate on the right are oriented as the boundaries of $\mathcal{W}$ and $\mathcal{B}_\rho$, respectively. But since $\operatorname{div}(\mathbf{F}) \equiv 0$ on $\mathcal{R}$, the triple integral on the left is of course 0, so the desired equality follows. □

(d) Use the previous parts (and some computation) to show that

$$\oiint_{\partial \mathcal{W}} \mathbf{F} \cdot d\mathbf{S} = 4\pi.$$

Solution. We've now reduced our original surface integral to an integral over $\partial \mathcal{B}_\rho$, which has unit normal vector $\mathbf{N}$ given by

$$\mathbf{N}(x,y,z) = \frac{1}{\sqrt{x^2+y^2+z^2}}\langle x,y,z\rangle.$$

Now

$$\oiint_{\partial \mathcal{W}} \mathbf{F} \cdot d\mathbf{S} = \oiint_{\partial \mathcal{B}_\rho} \mathbf{F} \cdot d\mathbf{S} = \oiint_{\partial \mathcal{B}_\rho} (\mathbf{F} \cdot \mathbf{N})\, dS,$$

so let's compute $\mathbf{F} \cdot \mathbf{N}$ on $\partial \mathcal{B}_\rho$. We have

$$\begin{aligned}\mathbf{F}(x,y,z) \cdot \mathbf{N}(x,y,z) &= \frac{1}{(x^2+y^2+z^2)^{3/2}} \frac{1}{\sqrt{x^2+y^2+z^2}} \langle x,y,z\rangle \cdot \langle x,y,z\rangle \\ &= \frac{x^2+y^2+z^2}{(x^2+y^2+z^2)^2} = \frac{1}{x^2+y^2+z^2} = \frac{1}{\rho^2},\end{aligned}$$

where the last equality uses the fact that $\partial \mathcal{B}_\rho$ is the sphere of radius ρ centered on the origin. So

$$\oiint_{\partial \mathcal{W}} \mathbf{F} \cdot d\mathbf{S} = \oiint_{\partial \mathcal{B}_\rho} \frac{1}{\rho^2}\, dS = \frac{1}{\rho^2}\,\text{area}(\partial \mathcal{B}_\rho) = \frac{4\pi\rho^2}{\rho^2} = \boxed{4\pi,}$$

as desired. □

(e) Explain why the following argument is wrong. First, $\text{div}(\mathbf{F}) \equiv 0$, so Poincaré's lemma says that $\mathbf{F}$ has a vector potential $\mathbf{G}$. So we can use Stokes' theorem:

$$\oiint_{\partial \mathcal{W}} \mathbf{F} \cdot d\mathbf{S} = \oiint_{\partial \mathcal{W}} \text{curl}(\mathbf{G}) \cdot d\mathbf{S} = \oint_{\partial(\partial \mathcal{W})} \mathbf{G} \cdot d\mathbf{r}.$$

But $\partial \mathcal{W}$ is a closed surface, so it doesn't have a boundary: $\partial(\partial \mathcal{W}) = \emptyset$. Therefore this last integral is 0, and we find that the surface integral in question evaluates to 0.

Solution. The error is in the very first claim. Namely, Poincaré's lemma only guarantees that divergence-free vector fields admit vector potentials on contractible domains. Since the domain of $\mathbf{F}$ is not contractible, we cannot use the fact that $\mathbf{F}$ is divergence-free to conclude that $\mathbf{F}$ has a vector potential. Indeed, the fact that the surface integral of $\mathbf{F}$ over the closed surface $\partial \mathcal{W}$ is nonzero demonstrates that $\mathbf{F}$ does not admit a vector potential. □

Such fun. Time for a more abstract take on all of this!

8.2 Day 21: Covectors and differential 1- and 2-forms

Goals

By the end of today's class, we should be able to do the following.

1. Define **covectors** and **differential 1-forms** on $\mathbb{R}^n$.
2. State the definition of the integral of a 1-form over a curve in $\mathbb{R}^n$.
3. Motivate the definition of **differential 2-forms** on $\mathbb{R}^n$.

As outlined above, we will now begin our attempt to generalize part two of the course to any dimension. The quick outline of part two is that we: introduced vector fields and their derivatives; defined integration

of vector fields over curves or surfaces; and finally expressed versions of the FTIC for these vector integrals. We've said before that vector fields can only be integrated over domains which have either dimension 1 or codimension 1, so we'll need a more complicated gadget than a vector field to serve as our integrand in this theory. Differential forms are precisely these gadgets.

A warning. It's difficult to give an intrinsic motivation for the material we're going to cover over the next few days. I'll do my best to give geometric interpretations for the objects we define, but we're also going to be developing a bit of algebraic machinery, and at times it may be difficult to match the algebraic maneuvers to their geometric meaning. You are strongly encouraged to read and reread this material, work the optional exercises, and ask questions. The final two weeks of the course will rely heavily on the ideas introduced in the next few days.

8.2.1 Covectors and differential 1-forms

We start by giving an abstract definition of covectors.

Definition. A **covector** on $\mathbb{R}^n$ is a linear map $\lambda\colon \mathbb{R}^n \to \mathbb{R}$. That is, λ is a map which satisfies

$$\lambda(\mathbf{v}+\mathbf{w}) = \lambda(\mathbf{v}) + \lambda(\mathbf{w}) \quad \text{and} \quad \lambda(c\mathbf{v}) = c\lambda(\mathbf{v}),$$

for any vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$ and any scalar $c \in \mathbb{R}$. We denote the collection of all covectors on $\mathbb{R}^n$ by $(\mathbb{R}^n)^*$.

With a sufficiently strong background in linear algebra, we can develop our entire theory of covectors from the above definition. But since we want to minimize linear algebra expectations, we'll take a slightly more grounded approach.

Example 8.1. We can define a covector dx on $\mathbb{R}^3$ which simply gives the x-coordinate of a vector:

$$dx\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = a.$$

You should verify that $dx\colon \mathbb{R}^3 \to \mathbb{R}$ is a linear map. We similarly define dy and dz.

Example 8.2. Generalizing Example 8.1 to any dimension, we can define a covector dx^k on $\mathbb{R}^n$ for each $1 \le k \le n$. In words, $dx^k(\mathbf{v})$ gives the k^{th} component of $\mathbf{v}$. As an equation, this is most easily written as a dot product:

$$dx^k(\mathbf{v}) := \mathbf{v} \cdot \mathbf{e}_k, \tag{8.1}$$

where $\mathbf{e}_k$ is the k^{th} standard basis vector for $\mathbb{R}^n$. We refer to the covectors $dx^1, \dots, dx^n$ as the **standard covectors** on $\mathbb{R}^n$.

Remark. The collection $(\mathbb{R}^n)^*$ of covectors on $\mathbb{R}^n$ is an n-dimensional vector space, called the *dual* of $\mathbb{R}^n$. The standard covectors $dx^1, \dots, dx^n$ form a basis for this dual space, and the standard covectors are by definition dual to the standard basis vectors of $\mathbb{R}^n$. The fact that we can relate these bases via equation 8.1 is sort of a special feature of the standard inner product on $\mathbb{R}^n$.

Example 8.3. Given any real numbers $c_1, \dots, c_n \in \mathbb{R}^n$, we can define a covector λ by

$$\lambda := c_1\, dx^1 + \cdots + c_n\, dx^n\colon \mathbb{R}^n \to \mathbb{R}.$$

So the standard covectors are said to *span* $(\mathbb{R}^n)^*$. For instance, consider the covector $\lambda = 3\,dx - 2\,dy + 4\,dz$. Then

$$\lambda\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = 3\,dx\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) - 2\,dy\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) + 4\,dz\left(\begin{bmatrix} a \\ b \\ c \end{bmatrix}\right) = 3a - 2b + 4c.$$

Exercise 8.4. Show that the standard covectors on $\mathbb{R}^n$ are **linearly independent**. That is, show that if

$$0 = c_1\, dx^1 + \cdots + c_n\, dx^n$$

for some real numbers $c_1, \dots, c_n$, then we must have $c_1 = \cdots = c_n = 0$.

What Example 8.3 and Exercise 8.4 show is that every covector on $\mathbb{R}^n$ can be uniquely written as a **linear combination** of the standard covectors. This motivates the following re-definition of covectors.

Definition. A **covector** on $\mathbb{R}^n$ is a linear combination of the standard covectors on $\mathbb{R}^n$. That is, a covector is a map $\lambda : \mathbb{R}^n \to \mathbb{R}$ of the form

$$\lambda = c_1\, dx^1 + \cdots + c_n\, dx^n,$$

for some real numbers $c_1, \ldots, c_n \in \mathbb{R}$.

This is all well and good, but what does a covector actually *do*? The answer is that a covector eats vectors, and outputs the component of the input vector along some fixed vector.

Example 8.5. Say we have a fixed vector $\mathbf{v}_0 := \langle a_0, b_0, c_0 \rangle \in \mathbb{R}^3$, and we want a map which take some vector $\mathbf{v} \in \mathbb{R}^3$ as input and tells us how much $\mathbf{v}$ pushes along $\mathbf{v}_0$. We can define a covector $\lambda = (\mathbf{v}_0)^*$ as

$$\lambda = a_0\, dx + b_0\, dy + c_0\, dz.$$

Then we notice that

$$\lambda(\mathbf{v}) = a_0(\mathbf{v} \cdot \mathbf{e}_1) + b_0(\mathbf{v} \cdot \mathbf{e}_2) + c_0(\mathbf{v} \cdot \mathbf{e}_3) = \mathbf{v} \cdot \mathbf{v}_0$$

gives us the component of $\mathbf{v}$ along $\mathbf{v}_0$. So *covectors are dual to vectors* (which is how they get the name).

Just as we can define vector fields, we can define covector fields. However, we will usually refer to covector fields as **differential 1-forms**. Actually, we'll often drop the word "differential" and just talk about **1-forms**.

Definition. A **differential 1-form** (or **1-form** or **covector field**) on $\mathbb{R}^n$ is a smooth[3]assignment of a covector on $\mathbb{R}^n$ to each point of $\mathbb{R}^n$. That is, a differential 1-form is a smooth map $\mathbb{R}^n \to (\mathbb{R}^n)^*$.

In practice, we won't think about the above definition too much. Instead, we'll consider a 1-form to be a linear combination of the **standard 1-forms** $dx^1, \ldots, dx^n$, except that scalar multiplication actually means multiplication by scalar-valued functions. That is, a differential 1-form on $\mathbb{R}^n$ can be written as

$$\lambda = F_1\, dx^1 + \cdots + F_n\, dx^n,$$

for some real-valued functions $F_1, \ldots, F_n : \mathbb{R}^n \to \mathbb{R}$.

Notation. It might be a little confusing that we used λ to denote a covector and also a covector field. Similarly, we used $dx^1, \ldots, dx^n$ to denote the standard covectors, and now they denote the standard 1-forms. This is because we can think of covectors as constant covector fields: they give us the same map at every point of $\mathbb{R}^n$. Going forward, we will think pretty much exclusively about differential 1-forms, and rarely about a single, stand-alone covector.

Just as we did with covectors, we give a more down-to-earth definition of differential 1-forms which is suited to the needs of this course. This is somehow less pleasing from an abstract point of view, but will meet our needs.

Definition. A **differential 1-form** on $\mathbb{R}^n$ is a linear combination over $C^\infty(\mathbb{R}^n)$ of the standard 1-forms on $\mathbb{R}^n$. That is, a differential 1-form λ is an expression

$$\lambda = F_1\, dx^1 + \cdots + F_n\, dx^n,$$

where $F_1, \ldots, F_n : \mathbb{R}^n \to \mathbb{R}$ are smooth maps. The set of all differential 1-forms on $\mathbb{R}^n$ is denoted $\Omega^1(\mathbb{R}^n)$.

Finally, we can define 1-forms on open subsets of $\mathbb{R}^n$ by restricting the above notation. A differential 1-form on an open subset $\mathcal{R} \subset \mathbb{R}^n$ is an expression of the form

$$F_1\, dx^1 + \cdots + F_n\, dx^n,$$

where $F_1, \ldots, F_n : \mathcal{R} \to \mathbb{R}$ are smooth maps. We denote the collection of all 1-forms on $\mathcal{R}$ by $\Omega^1(\mathcal{R})$.

[3]Making this precise would require thinking about a differential structure on the dual space $(\mathbb{R}^n)^*$. This isn't hard, but we won't do it. Instead, we'll think about 1-forms as linear combinations of standard 1-forms, where the "scalars" in our linear combinations are smooth functions $f : \mathbb{R}^n \to \mathbb{R}$.

Remark. Unlike $(\mathbb{R}^n)^*$, $\Omega^1(\mathbb{R}^n)$ is not a vector space. Instead, $\Omega^1(\mathbb{R}^n)$ is a *module* over the ring $C^\infty(\mathbb{R}^n)$. These notions are well beyond the scope of this course, but the basic idea is that, where scalar multiplication in $(\mathbb{R}^n)^*$ can be undone (provided the scalar is nonzero), the same is not true for $\Omega^1(\mathbb{R}^n)$. For instance, $\lambda = x\,dx$ is a differential 1-form on $\mathbb{R}$, and there does not exist any smooth function $F\colon \mathbb{R} \to \mathbb{R}$ with the property that $F\lambda = dx$.

8.2.2 Integrating 1-forms

There's enough abstraction going on that it might be difficult to keep track of what we're doing. But the mantra so far is that *covectors measure (signed) lengths of projections*. That is, every covector λ corresponds to some fixed vector $\mathbf{v}_0$, and evaluating λ on some vector $\mathbf{v}$ tells us the component of $\mathbf{v}$ along $\mathbf{v}_0$.

A mantra that many people find helpful for 1-forms is to say that *1-forms are objects which we can integrate over curves*. Indeed, if a covector tells us the component of a vector along some other vector, then we can use a covector field (i.e., a 1-form) to determine the component of the unit tangent vector $\mathbf{T}$ of some curve along a vector field. That's a lot of words. Here's an equation that might clear things up:

$$\int_{\mathcal{C}} \lambda := \int_{\mathcal{C}} \lambda(\mathbf{T})\,ds. \tag{8.2}$$

On the left, we have some 1-form λ which we want to integrate over an oriented curve $\mathcal{C}$. On the right, we have a scalar line integral over $\mathcal{C}$, and the scalar-valued function we're integrating is given by $\lambda(\mathbf{T})$.

Example 8.6. Say we have a pair of smooth functions $F_1, F_2\colon \mathbb{R}^2 \to \mathbb{R}$. Then we can define a differential 1-form $\lambda = F_1\,dx + F_2\,dy$ and consider the integral

$$\int_{\mathcal{C}} \lambda = \int_{\mathcal{C}} F_1\,dx + F_2\,dy,$$

for some reasonably nice curve $\mathcal{C}$. Now

$$\lambda(\mathbf{T}) = (F_1\,dx + F_2\,dy)(\mathbf{T}) = F_1(\mathbf{T}\cdot\mathbf{e}_1) + F_2(\mathbf{T}\cdot\mathbf{e}_2) = \langle F_1, F_2\rangle \cdot \mathbf{T}.$$

But this tells us that

$$\int_{\mathcal{C}} F_1\,dx + F_2\,dy = \int_{\mathcal{C}} \langle F_1, F_2\rangle \cdot \mathbf{T}\,ds = \int_{\mathcal{C}} \langle F_1, F_2\rangle \cdot d\mathbf{r},$$

where $\mathbf{F} = \langle F_1, F_2\rangle$. So the physicists' notation for line integrals makes sense!

At this point, it's totally reasonable if covectors and differential 1-forms seem like a big waste of time to you. In $\mathbb{R}^n$, every differential 1-form corresponds to some vector field, and integrating the 1-form over a curve is the same thing as integrating the vector field over the same curve. So why bother defining 1-forms? In a much larger context, there are a few reasons:

(a) integrating vector fields requires a choice of inner product — something we don't always have;

(b) by using the more abstract definitions (without reference to the standard covectors), we can develop a coordinate-free theory of integration;

(c) differential 1-forms set the stage for differential k-forms, which allow us to define integration over k-dimensional objects living in any ambient dimension.

We'll focus exclusively on reason (c). For instance, recall that we only defined surface integrals for surfaces that live in $\mathbb{R}^3$. In order to define integration on a surface living in $\mathbb{R}^n$, $n \geq 4$, we need the theory of 2-forms.

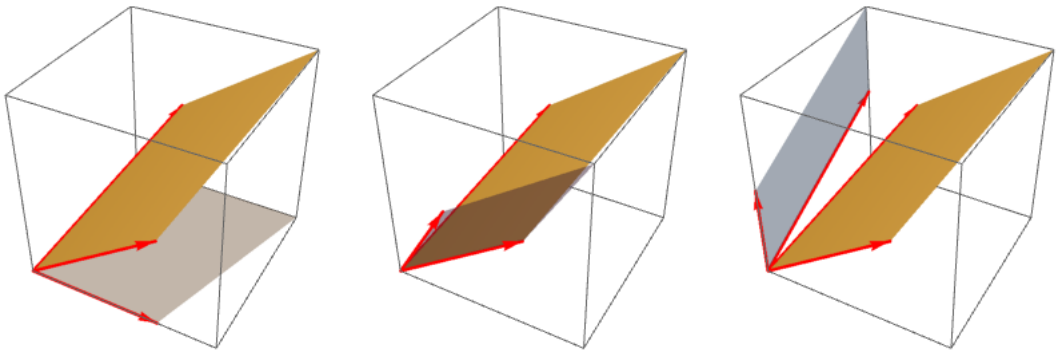

Figure 8.1: A parallelogram in $\mathbb{R}^3_{x,y,z}$ spanned by vectors $\mathbf{v}$ and $\mathbf{w}$, and its projections to the xy-, xz-, and yz-planes.

8.2.3 2-forms

Since we integrate 1-forms over curves, it seems natural that we should integrate 2-forms over surfaces. Unlike 1-forms, we won't give the abstract definition of differential 2-forms — it's just a bit *too* abstract. Instead, we'll define the **standard 2-forms** on $\mathbb{R}^n$ and, as we did with 1-forms, consider an arbitrary 2-form to be a linear combination of these.

We'll start our story in $\mathbb{R}^3$ for concreteness. Recall that the 1-forms dx, dy, and dz tell us the components of a vector $\mathbf{v}$ in the x-, y-, and z-directions, respectively. We want to similarly define standard 2-forms

$$dx \wedge dy, \quad dx \wedge dz, \quad \text{and} \quad dy \wedge dz$$

on $\mathbb{R}^3$. A 2-form will accept as input a pair $\mathbf{v}$, $\mathbf{w}$ of vectors and output a real number. In each case, this real number will correspond to the signed area of a projection of the parallelogram in $\mathbb{R}^3$ spanned by $\mathbf{v}$ and $\mathbf{w}$. Before giving a formulaic definition, we present an example.

Example 8.7. Consider the vectors $\mathbf{v} = \langle 7, 0, 5\rangle$ and $\mathbf{w} = \langle 3, 10, 5\rangle$ in $\mathbb{R}^3$. These vectors span a parallelogram in $\mathbb{R}^3$, and in Figure 8.1 we depict three different projections of this parallelogram: first onto the xy-plane, then onto the xz-plane, and finally onto the yz-plane. The signed areas of these three parallelograms should be given by

$$dx \wedge dy(\mathbf{v}, \mathbf{w}), \quad dx \wedge dz(\mathbf{v}, \mathbf{w}), \quad \text{and} \quad dy \wedge dz(\mathbf{v}, \mathbf{w}),$$

respectively. Let's focus on $dx \wedge dy(\mathbf{v}, \mathbf{w})$, which gives the signed area of the projection of our parallelogram to the xy-plane. This parallelogram is spanned by the projections of $\mathbf{v}$ and $\mathbf{w}$ to the xy-plane, which are given by

$$\text{proj}_{xy}(\mathbf{v}) = \langle 7, 0, 0\rangle \quad \text{and} \quad \text{proj}_{xy}(\mathbf{w}) = \langle 3, 10, 0\rangle.$$

When we say that we want the *signed* area, we mean that $dx \wedge dy(\mathbf{v}, \mathbf{w})$ should be positive if the vectors $\text{proj}_{xy}(\mathbf{v})$ and $\text{proj}_{xy}(\mathbf{w})$ are oriented in the same manner as $\partial_x = \mathbf{i}$ and $\partial_y = \mathbf{j}$. So we're computing the signed area of a parallelogram in $\mathbb{R}^2$. Indeed, since the z-coordinates of $\text{proj}_{xy}(\mathbf{v})$ and $\text{proj}_{xy}(\mathbf{w})$ are both zero, we can compute this area using a 2×2 determinant[4]

$$dx \wedge dy(\mathbf{v}, \mathbf{w}) := \det\big(\text{proj}_{xy}(\mathbf{v}) \quad \text{proj}_{xy}(\mathbf{w})\big) = \det\begin{pmatrix} 7 & 3 \\ 0 & 10 \end{pmatrix} = 70.$$

So if $\mathcal{P}$ is the parallelogram in $\mathbb{R}^3$ spanned by $\mathbf{v}$ and $\mathbf{w}$, then the signed area of the projection of $\mathcal{P}$ to the xy-plane is 70. Notice that swapping $\mathbf{v}$ and $\mathbf{w}$ multiplies this value by -1, since we have changed the orientation of $\mathcal{P}$:

$$dx \wedge dy(\mathbf{w}, \mathbf{v}) := \det\big(\text{proj}_{xy}(\mathbf{w}) \quad \text{proj}_{xy}(\mathbf{v})\big) = \det\begin{pmatrix} 3 & 7 \\ 10 & 0 \end{pmatrix} = -70.$$

[4]Of course we're abusing notation a bit here, since both $\text{proj}_{xy}(\mathbf{v})$ and $\text{proj}_{xy}(\mathbf{w})$ are vectors in $\mathbb{R}^3$ rather than $\mathbb{R}^2$. But we suppress the z-component, since this is forced to be 0.

We can obtain a similar sign flip by swapping the roles of dx and dy, but let's not explain this just yet. Instead, we point out that $dx \wedge dz$ and $dy \wedge dz$ are computed analogously: we project $\mathbf{v}$ and $\mathbf{w}$ to the relevant plane and then compute a relevant determinant. We have

$$dx \wedge dz(\mathbf{v}, \mathbf{w}) := \det\begin{pmatrix} \text{proj}_{xz}(\mathbf{v}) & \text{proj}_{xz}(\mathbf{w}) \end{pmatrix} = \det\begin{pmatrix} 7 & 3 \\ 5 & 5 \end{pmatrix} = 20$$

and

$$dy \wedge dz(\mathbf{v}, \mathbf{w}) := \det\begin{pmatrix} \text{proj}_{yz}(\mathbf{v}) & \text{proj}_{yz}(\mathbf{w}) \end{pmatrix} = \det\begin{pmatrix} 0 & 10 \\ 5 & 5 \end{pmatrix} = -50.$$

Okay, so the rule for computing the standard 2-forms in $\mathbb{R}^3$ can roughly be summarized as: delete the component we're ignoring, and compute the 2×2 determinant of what's left over. In order to generalize to higher dimensions, we'll need something a bit more formulaic.

Definition. Let $dx^1, \cdots, dx^n$ denote the standard differential 1-forms on $\mathbb{R}^n$. For every $1 \leq i, j \leq n$, define a 2-form $dx^i \wedge dx^j : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ by the formula

$$(dx^i \wedge dx^j)(\mathbf{v}, \mathbf{w}) := \det\left(\begin{pmatrix} v_1 & w_1 \\ \vdots & \vdots \\ v_n & w_n \end{pmatrix}_{(i,j)}\right),$$

for any vectors $\mathbf{v} = \langle v_1, \ldots, v_n \rangle, \mathbf{w} = \langle w_1, \ldots, w_n \rangle \in \mathbb{R}^n$. Here if M is an $n \times 2$ matrix, then $M_{(i,j)}$ is the 2×2 matrix whose first row is the i^{th} row of M and whose second row is the j^{th} row of M. The **standard 2-forms** on $\mathbb{R}^n$ are the 2-forms $dx^i \wedge dx^j$, with $i < j$.

Example 8.8. Consider the vectors

$$\mathbf{v} = \langle v_1, v_2, v_3, v_4 \rangle \quad \text{and} \quad \mathbf{w} = \langle w_1, w_2, w_3, w_4 \rangle$$

in $\mathbb{R}^4$. Let's compute $dx^3 \wedge dx^1(\mathbf{v}, \mathbf{w})$. Per the definition, this value is given by

$$dx^3 \wedge dx^1(\mathbf{v}, \mathbf{w}) = \det\left(\begin{pmatrix} v_1 & w_1 \\ v_2 & w_2 \\ v_3 & w_3 \\ v_4 & w_4 \end{pmatrix}_{(3,1)}\right).$$

Taking the third and first rows of this matrix as prescribed, we find that

$$dx^3 \wedge dx^1(\mathbf{v}, \mathbf{w}) = \det\begin{pmatrix} v_3 & w_3 \\ v_1 & w_1 \end{pmatrix} = v_3 w_1 - w_3 v_1.$$

Notice that we can also write this as

$$dx^3 \wedge dx^1(\mathbf{v}, \mathbf{w}) = dx^3(\mathbf{v}) dx^1(\mathbf{w}) - dx^3(\mathbf{w}) dx^1(\mathbf{v}).$$

This last expression is often used to define the standard 2-forms $dx^i \wedge dx^j$, but we'll find the determinant formulation somewhat more amenable to higher-order forms.

Exercise 8.9. Show that for any $1 \leq i, j \leq n$, we have $dx^j \wedge dx^i = -dx^i \wedge dx^j$. Use this to conclude that $dx^i \wedge dx^i = 0$, for any $1 \leq i \leq n$.

With the standard 2-forms in hand, we can state our definition of 2-forms on open subsets of $\mathbb{R}^n$.

Definition. A **differential 2-form** on an open subset $\mathcal{R}$ of $\mathbb{R}^n$ is a linear combination over $C^\infty(\mathcal{R})$ of the standard 2-forms on $\mathbb{R}^n$. That is, a differential 2-form ω is an expression

$$\omega = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} F_{i,j}\, dx^i \wedge dx^j,$$

where each $F_{i,j} : \mathcal{R} \to \mathbb{R}$ is a smooth map. The collection of all 2-forms on $\mathcal{R}$ is denoted by $\Omega^2(\mathcal{R})$.

Remark.

1. Notice that the summation in the definition of 2-forms is chosen to force $i < j$. This is because the standard 2-forms have $i < j$.
2. There are $\binom{n}{2} = \frac{1}{2}n(n-1)$ standard 2-forms on $\mathbb{R}^n$, and thus an arbitrary 2-form on $\mathcal{R}$ is determined by a choice of $\frac{1}{2}n(n-1)$ smooth functions on $\mathcal{R}$. Unless $n = 3$, this means that a 2-form does not correspond to a vector field on $\mathcal{R}$, unlike a 1-form.

8.2.4 Exercises

Exercise 8.10. For any vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^3$, verify the following formula for the cross product:

$$\mathbf{v} \times \mathbf{w} = \langle dy \wedge dz(\mathbf{v}, \mathbf{w}), dz \wedge dx(\mathbf{v}, \mathbf{w}), dx \wedge dy(\mathbf{v}, \mathbf{w}) \rangle .$$

Notice that the middle form $dz \wedge dx$ is not standard, according to our definition.

Exercise 8.11. Show that if $\phi : \mathbb{R}^n \to \mathbb{R}$ is any smooth function on $\mathbb{R}^n$ and $\mathbf{v} \in \mathbb{R}^n$ is an arbitrary vector, then

$$\nabla\phi \cdot \mathbf{v} = \frac{\partial \phi}{\partial x_1} dx^1(\mathbf{v}) + \cdots + \frac{\partial \phi}{\partial x_n} dx^n(\mathbf{v}).$$

The 1-form

$$d\phi := \frac{\partial \phi}{\partial x_1} dx^1 + \cdots + \frac{\partial \phi}{\partial x_n} dx^n$$

is often called the **total differential** of ϕ.

Exercise 8.12. We have established a correspondence between 1-forms on $\mathbb{R}^n$ and vector fields on $\mathbb{R}^n$. Say we have a 1-form $\lambda = F_1\, dx + F_2\, dy + F_3\, dz$ on $\mathbb{R}^3$, corresponding to the vector field $\mathbf{F} = \langle F_1, F_2, F_3 \rangle$. Now consider the 2-form $d\lambda$ defined by

$$d\lambda := dF_1 \wedge dx + dF_2 \wedge dy + dF_3 \wedge dz.$$

Here the symbol $d\lambda$ is (for now) meaningless, except as we've just defined it. Each of dF_1, dF_2, and dF_3 is a total differential, as defined in Exercise 8.11, and we compute, for instance, $dF_1 \wedge dx$, by "distributing the $\wedge$" as expected[5]. Show that

$$d\lambda = \mathrm{curl}_x(\mathbf{F})\, dy \wedge dz + \mathrm{curl}_y(\mathbf{F})\, dz \wedge dx + \mathrm{curl}_z(\mathbf{F})\, dx \wedge dy,$$

where $\mathrm{curl}_x(\mathbf{F})$ denotes the x-component of the vector field curl($\mathbf{F}$), and similarly for $\mathrm{curl}_y(\mathbf{F})$ and $\mathrm{curl}_z(\mathbf{F})$.

Exercise 8.13. (**Challenge**) Consider a smooth function $f : \mathbb{R}^n \to \mathbb{R}$, and let $\mathcal{S} \subset \mathbb{R}^n$ be the hypersurface defined by

$$\mathcal{S} := \{(x_1, \ldots, x_n) : f(x_1, \ldots, x_n) = c\},$$

for some $c \in \mathbb{R}$. Assume that for every point $p \in \mathcal{S}$, the derivatives $f_{x_1}(p), \ldots, f_{x_n}(p)$ do not all vanish. (This ensures that $\mathcal{S}$ is a *smooth* surface.) Prove that if $\mathbf{v} \in \mathbb{R}^n$ is tangent to $\mathcal{S}$ at a point $p \in \mathcal{S}$, then $df(\mathbf{v}) = 0$. Here df is the total differential of f.
Hint: Use the result from Exercise 8.11.

[5] We'll define the *wedge product* in our next class meeting. You should go ahead and attempt this problem in order to get comfortable with the notation.

8.3 Day 22: Differential k-forms and algebraic operations

Goals

By the end of today's class, we should be able to do the following.

1. Define **differential k-forms**.
2. Compute the **wedge product** of a pair of differential forms on $\mathbb{R}^n$.

8.3.1 Differential k-forms

In our last meeting we defined **1-forms** and **2-forms**, which measure signed lengths and areas[6], respectively. Namely, a 1-form λ on $\mathbb{R}^n$ corresponds to some vector field $\mathbf{F}$ on $\mathbb{R}^n$, and the value $\lambda(\mathbf{v})$ measures the component of a vector $\mathbf{v}$ along $\mathbf{F}$. Similarly, a 2-form ω ought to[7] correspond to some plane field[8] P, and $\omega(\mathbf{v},\mathbf{w})$ should[9] measure the signed area of the parallelogram spanned by $\mathbf{v}$ and $\mathbf{w}$, when projected to P.

By analogy, a k-form on $\mathbb{R}^n$ should[10] measure k-dimensional signed volumes. Namely, if η is a k-form on $\mathbb{R}^n$, then at every point $p \in \mathbb{R}^n$ we should have a map

$$\eta_p \colon \overbrace{\mathbb{R}^n \times \cdots \times \mathbb{R}^n}^{k \text{ times}} \to \mathbb{R},$$

and we imagine[11] that this map corresponds to an oriented k-dimensional subspace $V \subset \mathbb{R}^n$ which passes through p. Then the quantity

$$\eta_p(\mathbf{v}_1, \ldots, \mathbf{v}_k)$$

ought[12] to measure the signed volume of the parallelotope spanned by $\mathbf{v}_1, \ldots, \mathbf{v}_k$ when projected to V. For instance, the differential 3-form $dx^{(1,2,3)}$ on $\mathbb{R}^n$, $n \geq 3$ should be computed as follows:

$$dx^{(1,2,3)}\left(\begin{bmatrix} u_1 \\ \vdots \\ u_n \end{bmatrix}, \begin{bmatrix} v_1 \\ \vdots \\ v_n \end{bmatrix}, \begin{bmatrix} w_1 \\ \vdots \\ w_n \end{bmatrix}\right) = \det \begin{pmatrix} u_1 & v_1 & w_1 \\ u_2 & v_2 & w_2 \\ u_3 & v_3 & w_3 \end{pmatrix}.$$

That is, we project the vectors $\mathbf{u}$, $\mathbf{v}$, and $\mathbf{w}$ to the $x_1x_2x_3$ 3-plane, and compute signed volume within that space. Notice that, according to our intuition, a k-form on $\mathbb{R}^n$ is only a reasonable thing to consider if $k \leq n$. If $k > n$, then there is just one k-form on $\mathbb{R}^n$: the zero form.

Our definition of differential k-forms will mimic the presentation we gave last time for 2-forms, using determinants. In order to state this definition, we'll need to establish some notation.

Definition. Let M be an $n \times m$ matrix, and let $I = (i_1, \ldots, i_k)$ be a **multi-index** of length k, with $1 \leq i_j \leq n$ for each $1 \leq j \leq k$. That is, I is a sequence of k integers, each between 1 and n, with repetition allowed. The $k \times m$ matrix M_I is defined so that the j^{th} row of M_I is given by the $i_j{}^{\text{th}}$ row of M.

Example 8.14. Let

$$M = \begin{pmatrix} 2 & 9 & 4 \\ 7 & 5 & 3 \\ 6 & 1 & 8 \\ a & b & c \end{pmatrix}.$$

[6] This is a lie. The *standard* 2-forms measure signed areas, but in dimensions greater than three, a linear combination of standard 2-forms need not measure signed areas. Nonetheless, we will use this as a crutch for our intuition.

[7] But needn't.

[8] That is, a choice of two-dimensional plane at every point in $\mathbb{R}^n$.

[9] Nope.

[10] Again, lies.

[11] In general, we imagine incorrectly.

[12] This one's in italics, so we'll let it sli–nope! Not all k-forms correspond to k-spaces.

Then

$$M_{(3,2,1)} = \begin{pmatrix} 6 & 1 & 8 \\ 7 & 5 & 3 \\ 2 & 9 & 4 \end{pmatrix}, \quad M_{(1,4)} = \begin{pmatrix} 2 & 9 & 4 \\ a & b & c \end{pmatrix}, \quad \text{and} \quad M_{(2,2,2,2,2)} = \begin{pmatrix} 7 & 5 & 3 \\ 7 & 5 & 3 \\ 7 & 5 & 3 \\ 7 & 5 & 3 \\ 7 & 5 & 3 \end{pmatrix}.$$

Definition. With this notation established, we can now define a k-form $dx^I := dx^{i_1} \wedge \cdots \wedge dx^{i_k}$ on $\mathbb{R}^n$ for each multi-index $I = (i_1, \ldots, i_k)$ of integers between 1 and n, provided $1 \leq k \leq n$. This form is a map

$$dx^I : \overbrace{\mathbb{R}^n \times \cdots \times \mathbb{R}^n}^{k \text{ times}} \to \mathbb{R}$$

defined by $dx^I(\mathbf{v}_1, \ldots, \mathbf{v}_k) := \det(M_I)$, where

$$M = \begin{pmatrix} | & & | \\ \mathbf{v}_1 & \cdots & \mathbf{v}_k \\ | & & | \end{pmatrix}$$

is the $n \times k$ matrix whose columns are given by the vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ in $\mathbb{R}^n$.

Remark. Depending on how much matrix algebra you've seen, you may not be comfortable computing determinants of $n \times n$ matrices when n is greater than 3. From a conceptual standpoint, the most important thing we can say at this stage is the following:

In any dimension, determinants compute signed volumes.

If M is an $n \times n$ matrix, then $\det M$ computes the signed volume of the n-dimensional parallelotope spanned by the columns of M. In the above formulation, where I is a multi-index of length k, the matrix M_I is of size $k \times k$, and thus dx^I computes k-dimensional volumes. A possibly-more-comforting second remark is that in this class you probably won't have to actually compute any determinants of size bigger than 3×3. Just in case this changes, a 4×4 determinant is computed below. But it's more important that you know some properties of the determinant.

In the course of your mathematical career, you'll need to learn many properties of the determinant. Some of the most relevant properties for us are the following. Throughout, M is a $k \times k$ matrix.

1. For any constant c, $\det(cM) = c^k \det(M)$.
2. Suppose $1 \leq j \leq k$, and the j^{th} column $\mathbf{c}_j$ of M can be written as $\mathbf{c}_j = \mathbf{v} + \mathbf{w}$, for some $k \times 1$ vectors $\mathbf{v}, \mathbf{w}$. Let $M_{\mathbf{v}}$ and $M_{\mathbf{w}}$ denote the matrices obtained by replacing $\mathbf{c}_j$ with $\mathbf{v}$ and with $\mathbf{w}$, respectively. Then
$$\det(M) = \det(M_{\mathbf{v}}) + \det(M_{\mathbf{w}}).$$
Similarly, if the j^{th} row $\mathbf{r}_j$ has the form $\mathbf{r}_j = \mathbf{v} + \mathbf{w}$, then
$$\det(M) = \det(M^{\mathbf{v}}) + \det(M^{\mathbf{w}}),$$
where $M^{\mathbf{v}}$ and $M^{\mathbf{w}}$ denote the matrices obtained by replacing the row $\mathbf{r}_j$ with $\mathbf{v}$ and with $\mathbf{w}$, respectively.
3. For some integers $1 \leq i \neq j \leq k$, let $\tilde{M}$ be the matrix obtained by swapping the i^{th} and j^{th} columns of M. Then $\det(\tilde{M}) = -\det(M)$. Similarly, swapping two rows of M will multiply the determinant by -1.
4. If a given row or column of M is scaled by some scalar $c \in \mathbb{R}$, then $\det(M)$ is scaled by c. (Property 1 can be obtained from this property.)
5. Suppose A is an $k \times k$ matrix and D is an $\ell \times \ell$ matrix. Then
$$\det \begin{pmatrix} A & 0 \\ 0 & D \end{pmatrix} = \det(A) \cdot \det(D).$$
The matrix whose determinant we compute here is of size $(k + \ell) \times (k + \ell)$, and is said to be a **block matrix**.

These properties may seem overwhelming at first, but you should verify that each of them fits with our geometric understanding of the determinant.

Exercise 8.15. Show that if a multi-index I has any repeated indices, then the corresponding differential form dx^I is 0.

Definition. If $I = (i_1, \ldots, i_k)$ is an **increasing multi-index** of integers between 1 and n, meaning that

$$i_1 < i_2 < \cdots < i_k,$$

then we call dx^I a **standard k-form**.

Remark. Choosing a standard k-form on $\mathbb{R}^n$ amounts to choosing k distinct integers from the set $\{1, \ldots, n\}$. Once the k integers are chosen, we obtain a standard k-form by constructing an increasing multi-index I and then considering dx^I. Thus we see that there are $\binom{n}{k} = \frac{n!}{k!(n-k)!}$ standard k-forms on $\mathbb{R}^n$.

Exercise 8.16. Given a not-necessarily-increasing, length k multi-index I of distinct integers between 1 and n and $k+1$ vectors $\mathbf{v}_1, \ldots, \mathbf{v}_{k+1} \in \mathbb{R}^n$, use determinant properties to rewrite the quantity

$$dx^I(\mathbf{v}_1, \ldots, \mathbf{v}_{k-1}, \mathbf{v}_k) + dx^I(\mathbf{v}_1, \ldots, \mathbf{v}_{k-1}, \mathbf{v}_{k+1})$$

as a single evaluation of a standard k-form. Notice that $\mathbf{v}_{k+1}$ is omitted from the first term and $\mathbf{v}_k$ is omitted from the second.

Definition. Given $1 \leq k \leq n$, a **differential k-form** (or **k-form**) η on an open subset $\mathcal{R}$ of $\mathbb{R}^n$ is a linear combination of the standard k-forms on $\mathcal{R}$. That is, η can be written

$$\eta = \sum_I F_I \, dx^I,$$

where the sum is taken over all increasing multi-indices of integers between 1 and n with length k, and each function $F_I : \mathcal{R} \to \mathbb{R}$ is smooth. For completeness, we define a **0-form** on $\mathcal{R}$ to be a smooth function $f : \mathbb{R}^n \to \mathbb{R}$, and declare that all k-forms are zero if $k > n$. The collection of all k-forms on $\mathbb{R}^n$ is denoted $\Omega^k(\mathbb{R}^n)$, and we say that a k-form has degree k, denoted $\deg(\eta) = k$.

Example 8.17. Consider the differential 4-form $\eta = x_4 \, dx^{(1,2,3,5)}$ on $\mathbb{R}^5$, and consider the vector fields

$$\mathbf{v}_1 = \mathbf{e}_1 + \mathbf{e}_2, \quad \mathbf{v}_2 = \mathbf{e}_2 + \mathbf{e}_3, \quad \mathbf{v}_3 = \mathbf{e}_3 + \mathbf{e}_4, \quad \mathbf{v}_4 = \mathbf{e}_4 + \mathbf{e}_5.$$

Let's compute $\eta(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4)$. We have

$$\begin{aligned}
\eta(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4) &= x_4 \, dx^{(1,2,3,5)} \left(\begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 1 \end{bmatrix} \right) \\
&= x_4 \det \left(\begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix}_{(1,2,3,5)} \right) \\
&= x_4 \det \begin{pmatrix} 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.
\end{aligned}$$

We can compute this last determinant by the method of **cofactor expansion**[13]:

$$\begin{aligned}\det\begin{pmatrix}1&0&0&0\\1&1&0&0\\0&1&1&0\\0&0&0&1\end{pmatrix} &= (1)\det\begin{pmatrix}1&0&0\\1&1&0\\0&0&1\end{pmatrix} - (0)\det\begin{pmatrix}1&0&0\\0&1&0\\0&0&1\end{pmatrix}\\ &\quad + (0)\det\begin{pmatrix}1&1&0\\0&1&0\\0&0&1\end{pmatrix} - (0)\det\begin{pmatrix}1&1&0\\0&1&1\\0&0&0\end{pmatrix}\\ &= (1)\left((1)\det\begin{pmatrix}1&0\\0&1\end{pmatrix} - (0)\det\begin{pmatrix}1&0\\0&1\end{pmatrix} + (0)\det\begin{pmatrix}1&1\\0&0\end{pmatrix}\right)\\ &= (1)((1)(1) - (0)(1) + (0)(0)) = 1.\end{aligned}$$

Altogether, we find that

$$\boxed{\eta(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4) = x_4.}$$

We define addition and scalar multiplication of k-forms in the obvious way. That is, k-forms are added according to the formula

$$\left(\sum_I f_I\, dx^I\right) + \left(\sum_I g_I\, dx^I\right) := \sum_I (f_I + g_I)\, dx^I,$$

and multiplication by a scalar function is defined by

$$g\left(\sum_I f_I\, dx^I\right) := \sum_I (g f_I)\, dx^I.$$

It is important to note that addition is only defined for forms of the same degree.

8.3.2 Wedge products

Now it's time to explain the funny symbol $\wedge$. We've said that a 1-form λ on $\mathbb{R}^n$ corresponds to some vector field $\mathbf{F}$ on $\mathbb{R}^n$, and for any vector $\mathbf{v}$, $\lambda(\mathbf{v})$ tells us the signed component of $\mathbf{v}$ along $\mathbf{F}$. Now suppose we have another 1-form η, corresponding to a vector field $\mathbf{G}$. If $\mathbf{F}$ and $\mathbf{G}$ are linearly independent, then $\mathbf{F}$ and $\mathbf{G}$ span a plane field P, and there is a 2-form which measures signed areas in P, using the oriented basis $(\mathbf{F}, \mathbf{G})$. This 2-form will be given by $\lambda \wedge \eta$.

For any $1 \le k, \ell \le n$, the wedge product will give a map

$$\wedge\colon \Omega^k(\mathbb{R}^n) \times \Omega^\ell(\mathbb{R}^n) \to \Omega^{k+\ell}(\mathbb{R}^n).$$

That is, the wedge product of a k-form and an ℓ-form is a $(k+\ell)$-form. (If $k + \ell > n$, this product is automatically zero.) Geometrically, this means that if[14] λ measures k-volumes and η measures ℓ-volumes, then $\lambda \wedge \eta$ measures $(k+\ell)$-volumes (provided the k-volumes measures by λ are independent of the ℓ-volumes measured by η). Making sense of this algebraically can be a bit of a mess, so we will define the wedge product on the differential forms dx^I, where I is a multi-index, and then extend to more general forms via the distributive property.

Definition. Let I and J be nonempty multi-indices of integers between 1 and n. The **wedge product** of the associated differential forms dx^I and dx^J on $\mathbb{R}^n$ is given by concatenation:

$$dx^I \wedge dx^J := dx^{I*J},$$

where

$$I * J := (i_1, \ldots, i_k, j_1, \ldots, j_\ell)$$

is the concatenation of the multi-indices $I = (i_1, \ldots, i_k)$ and $J = (j_1, \ldots, j_\ell)$.

[13] Since we have actual numbers here — rather than variables — an even better method is *asking a computer*.

[14] As we've (foot)noted previously, our intuition is guided by the idea that all k-forms measure k-volumes, but this isn't actually true.

Exercise 8.18. Check that the wedge product is associative. That is, show that

$$dx^I \wedge (dx^J \wedge dx^K) = (dx^I \wedge dx^J) \wedge dx^K,$$

for any nonempty multi-indices I, J, and K.

Exercise 8.19. Show that if $I = (i_1, \dots, i_k)$, then

$$dx^I = dx^{i_1} \wedge \cdots \wedge dx^{i_k}.$$

Going forward, we will more often use the notation on the right than that on the left.

Exercise 8.20. Verify the following skew-commutativity property for the wedge product: for any multi-indices I and J,

$$dx^J \wedge dx^I = (-1)^{|I||J|}\, dx^I \wedge dx^J,$$

where $|I|$ and $|J|$ denote the lengths of the multi-indices I and J, respectively.

Even if this definition doesn't feel too inspired, let's observe that it at least passes some sort of dimension test: in the above, dx^I measures k-volumes, dx^J measures ℓ-volumes, and $dx^I \wedge dx^J$ measures $(k+\ell)$-volumes.

We can now extend the wedge product to arbitrary differential forms by assuming that this product behaves as we would expect (with the exception of skew-commutativity).

Definition. Let $\mathcal{R}$ be an open subset of $\mathbb{R}^n$, and let $\Omega^\bullet(\mathcal{R})$ denote the set of all[15] differential forms on $\mathbb{R}^n$ (of any degree). The **wedge product**

$$\wedge : \Omega^\bullet(\mathcal{R}) \times \Omega^\bullet(\mathcal{R}) \to \Omega^\bullet(\mathcal{R})$$

is defined by the following conditions:

(1) for 0-forms, the wedge product is scalar multiplication: $f \wedge \omega = f\,\omega$, for any 0-form f and any $\omega \in \Omega^\bullet(\mathcal{R})$;

(2) $dx^I \wedge dx^J$ follows the previous definition, for any multi-indices I, J of integers between 1 and n;

(3) for any forms $\lambda, \eta \in \Omega^\bullet(\mathcal{R})$ and any smooth function f on $\mathcal{R}$, $(f\,\lambda) \wedge \eta = f\,(\lambda \wedge \eta)$;

(4) for any forms $\lambda, \eta, \omega \in \Omega^\bullet(\mathcal{R})$, the wedge product distributes according to

$$(\lambda + \eta) \wedge \omega = \lambda \wedge \omega + \eta \wedge \omega \quad \text{and} \quad \omega \wedge (\lambda + \eta) = \omega \wedge \lambda + \omega \wedge \eta,$$

assuming that λ and η have the same degree.

Example 8.21. Let's show that the wedge product of any 1-form with itself is 0. We start by writing our 1-form as

$$\lambda = F_1\, dx^1 + \cdots + F_n\, dx^n,$$

for some smooth functions $F_1, \dots, F_n$ on an open subset $\mathcal{R}$ of $\mathbb{R}^n$. Using the distributive law for wedge products, we find that

$$\lambda \wedge \lambda = \sum_{i=1}^{n} \sum_{j=1}^{n} F_i F_j\, dx^i \wedge dx^j.$$

From here it's not hard to reason that $\lambda \wedge \lambda$ vanishes. Indeed, $dx^i \wedge dx^i = 0$, and whenever $i \neq j$ we have $dx^j \wedge dx^i = -dx^i \wedge dx^j$. So the addition of the term $F_j F_i\, dx^j \wedge dx^i$ corresponds to subtracting the term $F_i F_j\, dx^i \wedge dx^j$, canceling out all terms of our sum. Somewhat more carefully, can rewrite our sum as

$$\lambda \wedge \lambda = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} (F_i F_j\, dx^i \wedge dx^j + F_j F_i\, dx^j \wedge dx^i) = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} (F_i F_j\, dx^i \wedge dx^j - F_j F_i\, dx^i \wedge dx^j) = 0.$$

Example 8.22. Show that if λ and η are 1-forms on an open subset $\mathcal{R}$ of $\mathbb{R}^n$, then

$$(\lambda \wedge \eta)(\mathbf{v}, \mathbf{w}) = \lambda(\mathbf{v})\eta(\mathbf{w}) - \lambda(\mathbf{w})\eta(\mathbf{v}),$$

for any vector fields $\mathbf{v}, \mathbf{w}$ on $\mathcal{R}$. Explain how the result of Example 8.21 can be obtained from this fact.
Hint: First prove this under the assumption that λ and η are standard 1-forms, then extend via linearity.

[15] This notation is potentially dangerous: in case you're tempted to think otherwise, we point out that an element of $\Omega^\bullet(\mathcal{R})$ is a k-form on $\mathcal{R}$, for some $k \geq 0$, *not* a sum of forms of various degree.

Exercise 8.23. Generalizing Exercise 8.20, show that for any forms α and β,

$$\boxed{\beta \wedge \alpha = (-1)^{\deg(\alpha)\deg(\beta)} \alpha \wedge \beta.}$$

Explain how the result of Example 8.21 can be derived from this formula.

Example 8.24. Forms of degree greater than 1 need not square to zero, however. Consider coordinates (x_1, x_2, y_1, y_2) on $\mathbb{R}^4$. We can define a 2-form

$$\omega := dx^1 \wedge dy^1 + dx^2 \wedge dy^2.$$

Then

$$\begin{aligned}
\omega \wedge \omega &= (dx^1 \wedge dy^1 + dx^2 \wedge dy^2) \wedge (dx^1 \wedge dy^1 + dx^2 \wedge dy^2) \\
&= (dx^1 \wedge dy^1) \wedge (dx^1 \wedge dy^1) + (dx^2 \wedge dy^2) \wedge (dx^1 \wedge dy^1) \\
&\qquad + (dx^1 \wedge dy^1) \wedge (dx^2 \wedge dy^2) + (dx^2 \wedge dy^2) \wedge (dx^2 \wedge dy^2) \\
&= 0 + 2(dx^1 \wedge dy^1) \wedge (dx^2 \wedge dy^2) + 0 \\
&= -2(dx^1 \wedge dx^2 \wedge dy^1 \wedge dy^2).
\end{aligned}$$

The penultimate inequality makes use of Exercises 8.15 and 8.20.

Examples 8.21 and 8.24 are related to our footnotes about the lie of pretending that every k-form measures signed k-volume. For $k = 1$, it really is true that every 1-form measures signed length. Thus it makes sense that if λ is a 1-form, then $\lambda \wedge \lambda$ should be zero: if λ is associated to $\mathbf{F}$, then $\lambda \wedge \lambda$ is a 2-form associated to the parallelogram spanned by $\mathbf{F}$ and $\mathbf{F}$. But this "parallelogram" is one-dimensional, so it has no area. On the other hand, the 2-form we give in Example 8.24 does not correspond to any parallelogram. If it did, then the same geometric reasoning would lead us to conclude that $\omega \wedge \omega = 0$.

Example 8.25. In $\mathbb{R}^n$, we can use wedge products to compute the volume of an n-dimensional parallelotope without reference to the inner product on $\mathbb{R}^n$ (but in a way that agrees with the volume given by the inner product). Namely, we can define the **volume form** $\nu := dx^1 \wedge \cdots \wedge dx^n$ on $\mathbb{R}^n$, and we find that

$$|\nu(\mathbf{v}_1, \ldots, \mathbf{v}_n)| = \text{vol}(V), \tag{8.3}$$

where $\text{vol}(V)$ gives the volume (as determined by the inner product structure) of the parallelotope V spanned by $\mathbf{v}_1, \ldots, \mathbf{v}_n$. In a setting more general than ours, we might have access to differential forms, but not to an inner product. In that setting, we can take (8.3) as the definition of n-dimensional volume: we choose a top-degree differential form on our space and declare it to be the volume form.

Exercise 8.26. We can verify equation (8.3) as follows. In any inner product space W (you may think of W as $\mathbb{R}^n$), we can define the **Gramian** or **Gram determinant** of a collection of vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k \in W$:

$$G(\mathbf{v}_1, \ldots, \mathbf{v}_k) := \det \begin{pmatrix} \mathbf{v}_1 \cdot \mathbf{v}_1 & \mathbf{v}_1 \cdot \mathbf{v}_2 & \cdots & \mathbf{v}_1 \cdot \mathbf{v}_k \\ \mathbf{v}_2 \cdot \mathbf{v}_1 & \mathbf{v}_2 \cdot \mathbf{v}_2 & \cdots & \mathbf{v}_2 \cdot \mathbf{v}_k \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{v}_k \cdot \mathbf{v}_1 & \mathbf{v}_k \cdot \mathbf{v}_2 & \cdots & \mathbf{v}_k \cdot \mathbf{v}_k \end{pmatrix}.$$

Now the vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ span a polytope V of dimension at most k, and the k-dimensional volume of this polytope is given (in the inner product space formulation) by

$$\text{vol}_k(V) = \sqrt{G(\mathbf{v}_1, \ldots, \mathbf{v}_k)}.$$

Taking this as the definition of volume, and with the volume form[16] on $\mathbb{R}^n$ defined as above, prove that

$$|\nu(\mathbf{v}_1, \ldots, \mathbf{v}_n)| = \text{vol}_n(V),$$

for any vectors $\mathbf{v}_1, \ldots, \mathbf{v}_n \in \mathbb{R}^n$.

[16] Let's be clear about what we're glossing over here. The form ν is our preferred volume form precisely because it gives us the same volume as does the inner product. Without the knowledge of the inner product structure on $\mathbb{R}^n$, we wouldn't necessarily have a reason to prefer this volume form over some other top-degree form.

Exercise 8.27. Given a choice of volume form ν on $\mathbb{R}^n$, we can compute n-volumes without reference to the inner product. But what about k-volumes, where $1 \leq k < n$? Say we have linearly independent vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ in $\mathbb{R}^n$, and let V be the k-dimensional parallelotope spanned by these vectors. We can't feed $\mathbf{v}_1, \ldots, \mathbf{v}_k$ into ν, so how do we determine the k-dimensional volume of V? Our basic intuition is as follows: we can find the area of a rectangle in $\mathbb{R}^3$ by turning the rectangle into a rectangular prism, where the "height" above the rectangle is 1. The area of the rectangle is then the same as the volume of the resulting rectangular prism. In our more general setting, we can choose an *orthonormal basis* for the *orthogonal complement* $V^\perp$ of V. That is, we choose vectors $\mathbf{n}_1, \ldots, \mathbf{n}_{n-k}$ such that

$$\mathbf{n}_i \cdot \mathbf{v}_j = 0 \quad \text{and} \quad \mathbf{n}_i \cdot \mathbf{n}_j = \begin{cases} 0, & i \neq j \\ 1, & i = j \end{cases}.$$

Notice that our ability to choose these vectors relies on the inner product structure of $\mathbb{R}^n$. Prove that

$$|\nu(\mathbf{v}_1, \ldots, \mathbf{v}_k, \mathbf{n}_1, \ldots, \mathbf{n}_{n-k})| = \mathrm{vol}_k(V),$$

where $\mathrm{vol}_k(V)$ is the k-dimensional volume of V, as defined by the Gramian.
Hint: Use Exercise 8.26 to rewrite the desired equality in terms of Gramians. You should end up needing to compute the determinant of a block matrix.

9 Week 9

This week we carry on in our quest to generalize part two of the course. We start by defining **differentiation** for differential forms — an important ingredient in any FTIC — as well as a notion of change-of-variables for forms. We then identify the subsets of $\mathbb{R}^n$ over which we will integrate k-forms, and finally define the integral of a k-form over a **parametrized k-surface**.

9.1 Day 23: Differentiation

Goals

By the end of today's class, we should be able to do the following.

1. Define and interpret the **exterior derivative** of a differential form.

9.1.1 Differentiation

We want to define a notion of differentiation for differential forms which at least moderately agrees with the philosophy that derivatives should measure how quickly a quantity is changing. The only possibly-artificial requirement we impose on our differentiation is that we want the derivative of a differential form to be another differential form.

0-forms

Let's start with the easiest case: 0-forms. A 0-form is just a smooth function $f : \mathbb{R}^n \to \mathbb{R}$, and we're quite familiar with derivatives for smooth functions. Our preferred derivative here will be the *directional derivative*. That is, the derivative of a 0-form f is the 1-form $df : \mathbb{R}^n \to \mathbb{R}$ defined by

$$df(\mathbf{v}) := D_{\mathbf{v}} f = \nabla f \cdot \mathbf{v}.$$

We've seen this 1-form before: according to Exercise 8.11, df is the **total differential** of f.

Definition. Let $f \in \Omega^0(\mathcal{R})$ be a 0-form on some open subset $\mathcal{R}$ of $\mathbb{R}^n$. The **derivative** (or **exterior derivative**) df of f is the total differential of f.

Exercise 9.1. Prove the product rule for total differentials: $d(fg) = g\,df + f\,dg$, for any smooth functions $f, g : \mathbb{R}^n \to \mathbb{R}$.

1-forms

That wasn't so bad. We started with a 0-form f and asked, given a vector $\mathbf{v}$, how quickly f changes as we move in the direction (and at the speed) determined by $\mathbf{v}$. Maybe we can do something similar for a 1-form $\lambda \in \Omega^1(\mathcal{R})$. Let's ask the vague question, "How does λ vary as we move in the direction of $\mathbf{v}$?"

The first problem we run into is nailing down what we mean by, "How does λ vary?" In order to get a real number out of λ, we need to plug in a vector — this is in addition to the vector $\mathbf{v}$ along which we're taking a derivative. So we realize that *the derivative of a 1-form will eat two vectors and output a real number*. Since we expect derivatives of differential forms to be differential forms, we conclude that the derivative of a 1-form will be a 2-form.

Remark. One point about which we could have been more careful is that a differential k-form on $\mathcal{R}$ doesn't just eat k vectors — it eats a point p in $\mathcal{R}$, and then k vectors based at p. A typical notation for this is

$$\eta_p(\mathbf{v}_1, \ldots, \mathbf{v}_k),$$

to indicate that the vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ are all based at p. So if $\mathbf{w}$ is a vector field on $\mathcal{R}$, then we can think of the quantity $\lambda(\mathbf{w})$ as a smooth function $\mathcal{R} \to \mathbb{R}$: for each point $p \in \mathcal{R}$, we compute $\lambda_p(\mathbf{w}(p))$. This perspective should be helpful as we define the derivative.

A reasonable first attempt at defining $d\lambda(\mathbf{v},\mathbf{w})$ would treat $\mathbf{w}$ as a (constant) vector field and simply compute the directional derivative $D_{\mathbf{v}}(\lambda(\mathbf{w}))$. Since $\lambda(\mathbf{w})$ is a smooth function, we can take its directional derivative in the direction of $\mathbf{v}$, so this makes sense. The problem is that $D_{\mathbf{v}}(\lambda(\mathbf{w}))$ treats the vectors $\mathbf{v}$ and $\mathbf{w}$ very unequally. We want $d\lambda$ to be a 2-form, but there's no reason to expect the expression $D_{\mathbf{v}}(\lambda(\mathbf{w}))$ to be alternating in the vectors $\mathbf{v}$ and $\mathbf{w}$. We fix this by taking the *alternatization* of the expression.

Definition. Let $\lambda \in \Omega^1(\mathcal{R})$ be a differential 1-form on some open subset $\mathcal{R}$ of $\mathbb{R}^n$. The **derivative** (or **exterior derivative**) $d\lambda$ of λ is a 2-form on $\mathcal{R}$ defined by

$$d\lambda(\mathbf{v},\mathbf{w}) := D_{\mathbf{v}}(\lambda(\mathbf{w})) - D_{\mathbf{w}}(\lambda(\mathbf{v})),$$

for vectors $\mathbf{v},\mathbf{w}$, treated as constant vector fields on $\mathcal{R}$.

Remark. In a more abstract setting, it's not as easy to just pretend that vectors are in fact vector fields; i.e., there's not generally such a thing as a constant vector field. Writing down this definition of the derivative in such a setting is a little more work.

Exercise 9.2. Show that exterior differentiation is a linear map $d\colon \Omega^1(\mathcal{R}) \to \Omega^2(\mathcal{R})$. That is, show that if $c \in \mathbb{R}$ is a scalar and λ, η are 1-forms on $\mathcal{R}$, then

$$d(c\,\lambda) = c\,d\lambda \quad \text{and} \quad d(\lambda+\eta) = d\lambda + d\eta.$$

Note: It's important that c is a constant scalar, and not a scalar function.

Example 9.3. Let's compute the derivative of a standard 1-form on $\mathbb{R}^n$. Take $1 \le j \le n$ and $\mathbf{v},\mathbf{w} \in \mathbb{R}^n$. Then

$$d(dx^j)(\mathbf{v},\mathbf{w}) = D_{\mathbf{v}}(dx^j(\mathbf{w})) - D_{\mathbf{w}}(dx^j(\mathbf{v})).$$

But moving our point in the direction of $\mathbf{v}$ doesn't affect the value of the function $dx^j(\mathbf{w})$, so $D_{\mathbf{v}}(dx^j(\mathbf{w})) = 0$. We similarly find that $D_{\mathbf{w}}(dx^j(\mathbf{v})) = 0$, so $\boxed{d(dx^j) = 0}$.

Example 9.4. Together, Exercise 9.2 and Example 9.3 might worry us. We've shown that the derivative of a standard 1-form is 0, and that differentiation is linear. Since an arbitrary 1-form is a linear combination (over $C^\infty(\mathcal{R})$) of standard 1-forms, we might be tempted to think that the derivative of a 1-form is always 0. Instead, we have a product rule. Suppose we have a 0-form f and a standard 1-form dx^j on $\mathcal{R} \subset \mathbb{R}^n$. Then for any vectors $\mathbf{v},\mathbf{w} \in \mathbb{R}^n$,

$$d(f\,dx^j)(\mathbf{v},\mathbf{w}) = D_{\mathbf{v}}(f\,dx^j(\mathbf{w})) - D_{\mathbf{w}}(f\,dx^j(\mathbf{v})).$$

But again, $dx^j(\mathbf{w})$ doesn't depend on $\mathbf{v}$, so our product rule for directional derivatives tells us that

$$\begin{aligned} d(f\,dx^j)(\mathbf{v},\mathbf{w}) &= D_{\mathbf{v}}(f)\,dx^j(\mathbf{w}) - D_{\mathbf{w}}(f)\,dx^j(\mathbf{v}) \\ &= df(\mathbf{v})\,dx^j(\mathbf{w}) - df(\mathbf{w})\,dx^j(\mathbf{v}) \\ &= (df \wedge dx^j)(\mathbf{v},\mathbf{w}). \end{aligned}$$

So we find that $\boxed{d(f\,dx^j) = df \wedge dx^j}$. The last equality depends on Exercise 8.22.

k-forms

Now that we've differentiated 1-forms, doing the same for k-forms shouldn't be too bad. First, if η is a k-form, then $d\eta$ is a $(k+1)$-form. As a naïve idea, this is because we feed k vectors (plus a point) into η in order to get a real number, then choose one more vector and ask how the real number changes as we move in this auxiliary direction. That is, the naïve (and incorrect) attempt at defining $d\eta$ would have us compute the directional derivative

$$D_{\mathbf{v}_{k+1}}(\eta(\mathbf{v}_1,\ldots,\mathbf{v}_k)),$$

where again we treat each vector $\mathbf{v}_i$ as a constant vector field. As before, this does not give equal treatment to each of the $k+1$ vectors. We put all of the vectors on equal footing by, one at a time, treating each as the vector which is varied while the others are plugged into η. We then sum the resulting derivatives, except that the requirement that $d\eta$ be a differential form (and thus skew-commutative) introduces a sign that we must track.

Definition. Let $\eta \in \Omega^k(\mathcal{R})$ be a differential k-form on some open subset $\mathcal{R}$ of $\mathbb{R}^n$. The **derivative** (or **exterior derivative**) $d\eta$ of η is a $(k+1)$-form on $\mathcal{R}$ defined by

$$d\eta(\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}) = \sum_{j=1}^{k+1} (-1)^{j+1} D_{\mathbf{v}_j}(\eta(\mathbf{v}_1, \ldots, \widehat{\mathbf{v}}_j, \ldots, \mathbf{v}_{k+1})),$$

for vectors $\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}$. Here the notation $\widehat{\mathbf{v}}_j$ indicates that we omit the vector $\mathbf{v}_j$, and all vectors are treated as constant vector fields.

Exercise 9.5. Verify that this definition of the derivative recovers the previous definitions for 0- and 1-forms.

Remark. In Exercise 9.2 we saw that the linearity of differentiation for 1-forms follows from the linearity of directional derivatives. The same reasoning tells us that differentiation of k-forms is linear, for any k.

Example 9.6. As was the case for 1-forms, the standard k-forms have derivative 0. Indeed, for any vectors $\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}$ and appropriate multi-index I,

$$d(dx^I)(\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}) = \sum_{j=1}^{k+1} (-1)^{j+1} D_{\mathbf{v}_j}(dx^I(\mathbf{v}_1, \ldots, \widehat{\mathbf{v}}_j, \ldots, \mathbf{v}_{k+1})).$$

But, just as we saw in the case $k = 1$, the quantity $dx^I(\mathbf{v}_1, \ldots, \widehat{\mathbf{v}}_j, \ldots, \mathbf{v}_{k+1})$ doesn't change as we vary our point in the direction of $\mathbf{v}_j$. So each of the directional derivatives in our sum vanishes, and we see that $\boxed{d(dx^I) = 0}$.

Exercise 9.7. Repeat the argument of Example 9.4 to show that $\boxed{d(f\,dx^I) = df \wedge dx^I}$ for any 0-form $f \in \Omega^0(\mathcal{R})$ and any standard k-form dx^I on $\mathcal{R}$.

Example 9.8. Let's prove a version of the product rule for differentiation of wedge products. This is often called a **Leibniz rule**. Say we have an open subset $\mathcal{R}$ of $\mathbb{R}^n$, along with a k-form α and an ℓ-form β on $\mathcal{R}$. Then we may write

$$\sum_I f_I\, dx^I \quad \text{and} \quad \sum_J g_J\, dx^J,$$

where the sums are taken over multi-indices of length k and ℓ, respectively. Using the algebraic properties of the wedge product, the linearity of differentiation, and the result of Exercise 9.7, we find that

$$\begin{aligned} d(\alpha \wedge \beta) &= d\left(\left(\sum_I f_I\, dx^I\right) \wedge \left(\sum_J g_J\, dx^J\right)\right) = d\left(\sum_{I,J} f_I g_J\, dx^I \wedge dx^J\right) \\ &= \sum_{I,J} d\left(f_I g_J\, dx^I \wedge dx^J\right) = \sum_{I,J} d(f_I g_J)\, dx^I \wedge dx^J. \end{aligned}$$

At this point we can apply the product rule derived in Exercise 9.1, as well as algebraic properties of wedge products, to find that

$$\begin{aligned} d(\alpha \wedge \beta) &= \sum_{I,J} g_J\, df_I \wedge dx^I \wedge dx^J + \sum_{I,J} f_I\, dg_J \wedge dx^I \wedge dx^J \\ &= \sum_{I,J} (df_I \wedge dx^I) \wedge (g_J\, dx^J) + (-1)^{\deg(dg_J)\deg(dx^I)} \sum_{I,J} (f_I\, dx^I) \wedge (dg_J \wedge dx^J) \\ &= d\alpha \wedge \beta + (-1)^k \alpha \wedge d\beta. \end{aligned}$$

So we have the general rule that

$$\boxed{d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^{\deg(\alpha)} \alpha \wedge d\beta}.$$

It's probably not obvious that the signs in this last computation are correct. You should check that they work out.

Exercise 9.9. Check the computations in Example 9.8 for sign errors.

The derivative allows us to introduce some terminology for k-forms. The main characters in part two of the course were vector fields, and we cared a lot about determining whether or not we could find a potential (function or vector field) for a given vector field. With differential forms, we can streamline this question a bit.

Definition. Let η be a k-form on an open subset $\mathcal{R}$ of $\mathbb{R}^n$. If $k \geq 1$, then we say that η is **exact** if there is a $(k-1)$-form λ satisfying $d\lambda = \eta$, in which case we call λ a **potential** for η. We say that η is **closed** if $d\eta = 0$.

Exercise 9.10. Identify all closed 0-forms.

In this new language, we would say that a conservative vector field is exact, since it is the derivative of some scalar potential function. Notions like curl-free are similarly encoded as statements about forms being closed. Recall that — by our definition of curl — conservative vector fields are curl-free. In the language of forms, this is reflected by the fact that *exact forms are closed*, which follows from the following exercise.

Exercise 9.11. Prove that $\boxed{d \circ d = 0}$ (often called the **nilpotence** of d). That is, show that if α is a k-form, $k \geq 0$, then $d(d\alpha) = 0$.
Hint: First prove this under the assumption that $\alpha = f\, dx^I$, and then extend by linearity.

Axioms for d. The derivative we've defined here can also be defined as the unique linear (over $\mathbb{R}$) map $d : \Omega^k(\mathcal{R}) \to \Omega^{k+1}(\mathcal{R})$ such that

(1) df is the total differential, for any 0-form f;

(2) $d(df) = 0$, for any 0-form f;

(3) $d(\alpha \wedge \beta) = d\alpha \wedge \beta + (-1)^{\deg(\alpha)}(\alpha \wedge d\beta)$, for any forms α and β.

Hopefully you're not too mad at me for holding off on these axioms until we'd seen the somewhat more elaborate definition. The axioms are useful for obtaining the relevant algebraic properties, but are maybe a bit harder to motivate geometrically.

Remember that we initially defined the vector derivatives curl and div algebraically, and delayed their geometric interpretations until we had established a relevant FTIC. A similar saga will play out here; after proving Stokes' theorem, we'll revisit the definition of d and give the "correct" geometric interpretation.

9.2 Day 24: Pullbacks

Goals

By the end of today's class, we should be able to do the following.

1. Compute the **pullback** of a differential form under a smooth map.

9.2.1 Pullback

Given a smooth map $\Phi : \mathcal{D} \to \mathcal{R}$ between open subsets $\mathcal{D} \subset \mathbb{R}^m$ and $\mathcal{R} \subset \mathbb{R}^n$, we want to know how Φ relates the linear structures of $\mathbb{R}^m$ and $\mathbb{R}^n$ which we see on $\mathcal{D}$ and $\mathcal{R}$. For instance, for any point $p \in \mathcal{D}$ and vector $\mathbf{v} \in \mathbb{R}^m$ based at p, we have a vector $\Phi_*\mathbf{v}$ in $\mathbb{R}^n$ based at $\Phi(p) \in \mathcal{R}$. Here Φ_* is the Jacobian matrix of Φ, and $\Phi_*\mathbf{v}$ is a matrix product, which we call the **pushforward** of $\mathbf{v}$. This tells us how to push a vector forward from $\mathcal{D}$ to $\mathcal{R}$. Similarly, we can define the **pullback** of a differential form.

Definition. Let $\Phi : \mathcal{D} \to \mathcal{R}$ be a smooth map between open subsets $\mathcal{D} \subset \mathbb{R}^m$ and $\mathcal{R} \subset \mathbb{R}^n$. The **pullback** of a 0-form $f \in \Omega^0(\mathcal{R})$ is the 0-form $\Phi^* f \in \Omega^0(\mathcal{D})$ defined by

$$\Phi^* f := f \circ \Phi.$$

For $k \geq 1$, the **pullback** of a k-form $\lambda \in \Omega^k(\mathcal{R})$ is the k-form $\Phi^*\lambda \in \Omega^k(\mathcal{D})$ defined by

$$(\Phi^*\lambda)_p(\mathbf{v}_1, \ldots, \mathbf{v}_k) := \lambda_{\Phi(p)}(\Phi_*\mathbf{v}_1, \ldots, \Phi_*\mathbf{v}_k),$$

for any vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k \in \mathbb{R}^m$ based at $p \in \mathcal{D}$, where Φ_* is the Jacobian matrix of Φ.

Remark. It's worth noting that it's much easier to pull a differential form back over a map than it is to push a vector field forward. We can always compute the pushforward of a single vector, but we cannot necessarily transform a vector field $\mathbf{F}$ on $\mathcal{D}$ into a vector field on $\mathcal{R}$ using a smooth map $\Phi\colon \mathcal{D} \to \mathcal{R}$. If Φ fails to be injective, then we can find a point $q \in \mathcal{R}$ with multiple preimages $p_1, p_2 \in \mathcal{D}$. How do we decide whether to push $\mathbf{F}(p_1)$ or $\mathbf{F}(p_2)$ forward? On the other hand, if Φ fails to be surjective, then we will find points in $\mathcal{R}$ at which there's no reasonable candidate for our pushforward vector. There are remedies to these problems, but you should convince yourself that pullbacks don't present the same sort of trouble.

Exercise 9.12. Verify that $\Phi^*(f\,\lambda) = (f \circ \Phi)\,\Phi^*\lambda$.

Example 9.13. As a first example, let's compute the pullback of a standard 1-form. We take $\Phi\colon \mathcal{D} \to \mathcal{R}$ as in the definition, and use coordinates $u_1, \dots, u_m$ on $\mathcal{D}$ and coordinates $x_1, \dots, x_n$ on $\mathcal{R}$. For some choice of smooth functions $f_{i1}, \dots, f_{im}$ on $\mathcal{D}$ we can write

$$\Phi^*(dx^i) = f_{i1}\,du^1 + \cdots + f_{im}\,du^m.$$

We can determine the function f_{ij} by evaluating $\Phi^*(dx^i)$ on the coordinate vector ∂_{u_j}. That is,

$$\Phi^*(dx^i)(\partial_{u_j}) = (f_{i1}\,du^1 + \cdots + f_{im}\,du^m)(\partial_{u_j}) = f_{ij},$$

since $du^j(\partial_{u_j}) = 1$, while the other standard 1-forms vanish on ∂_{u_j}. Now recall that, for $1 \le i \le n$ and any vector $\mathbf{v} \in \mathbb{R}^n$,

$$dx^i(\mathbf{v}) = \partial_{x_i} \cdot \mathbf{v}.$$

So we find that, for any $1 \le i \le n$ and $1 \le j \le m$,

$$\Phi^*(dx^i)(\partial_{u_j}) = dx^i(\Phi_*\partial_{u_j}) = \partial_{x_i} \cdot \frac{\partial \Phi}{\partial u_j} = \frac{\partial \Phi_i}{\partial u_j}.$$

The upshot is that $f_{ij} = \frac{\partial \Phi_i}{\partial u_j}$, so

$$\Phi^*(dx^i) = \frac{\partial \Phi_i}{\partial u_1}\,du^1 + \cdots + \frac{\partial \Phi_i}{\partial u_m}\,du^m, \tag{9.1}$$

for any $1 \le i \le n$. Notice that if we treat $\Phi_i\colon \mathcal{D} \to \mathbb{R}$ as a function in its own right, then (9.1) says

$$\Phi^*(dx^i) = d\Phi_i.$$

That is, the pullback of the i^{th} standard 1-form on $\mathcal{R}$ is the total differential of the i^{th} component of Φ.

Remark. The above example uses the notation ∂_{u_j} for a coordinate vector field. You've probably seen the standard basis vectors for $\mathbb{R}^n$ denoted by $\mathbf{e}_1, \dots, \mathbf{e}_n$, and also treated these vectors as fields. Because we'll be dealing with multiple Euclidean spaces at one time (for instance, $\mathbb{R}^m$ and $\mathbb{R}^n$), we will instead use the notation ∂_{u_j}. If $\mathbb{R}^m$ has coordinates $u_1, \dots, u_m$, then we will denote the standard basis vectors on $\mathbb{R}^m$ by $\partial_{u_1}, \dots, \partial_{u_m}$. Similarly, if we use coordinates $x_1, \dots, x_n$ on $\mathbb{R}^n$, then the standard basis vectors on $\mathbb{R}^n$ are denoted $\partial_{x_1}, \dots, \partial_{x_n}$. Going forward, we will use this notation for any coordinate system without comment.

Example 9.14. The pullback might seem intimidating at first, but it's really supposed to be automating some familiar change-of-coordinates ideas. For instance, consider cylindrical coordinates[1] $\Phi\colon \mathbb{R}^3_{r,\theta,z} \to \mathbb{R}^3_{x,y,z}$ defined by

$$\Phi(r, \theta, z) = (r\cos\theta, r\sin\theta, z).$$

We can use (9.1) to compute the pullbacks $\Phi^*(dx)$, $\Phi^*(dy)$, and $\Phi^*(dz)$. We have

$$\begin{aligned}
\Phi^*(dx) &= \frac{\partial}{\partial r}(r\cos\theta)\,dr + \frac{\partial}{\partial \theta}(r\cos\theta)\,d\theta + \frac{\partial}{\partial z}(r\cos\theta)\,dz = \cos\theta\,dr - r\sin\theta\,d\theta \\
\Phi^*(dy) &= \frac{\partial}{\partial r}(r\sin\theta)\,dr + \frac{\partial}{\partial \theta}(r\sin\theta)\,d\theta + \frac{\partial}{\partial z}(r\sin\theta)\,dz = \sin\theta\,dr + r\cos\theta\,d\theta \\
\Phi^*(dz) &= \frac{\partial}{\partial r}(z)\,dr + \frac{\partial}{\partial \theta}(z)\,d\theta + \frac{\partial}{\partial z}(z)\,dz = dz.
\end{aligned}$$

In our less rigorous days we might have written

$$\begin{aligned}
x &= r\cos\theta &&\Rightarrow && dx &= \cos\theta\,dr - r\sin\theta\,d\theta \\
y &= r\sin\theta &&\Rightarrow && dy &= \sin\theta\,dr + r\cos\theta\,d\theta \\
z &= z &&\Rightarrow && dz &= dz,
\end{aligned}$$

[1] As written here, Φ is not injective, but the pullback doesn't care about this.

where the equations on the right are obtained as a sort of total differential. Now that we know what the symbols dx, dy, and dz mean, we can express this change of variables more accurately. In practice, our scratch computations are done just as before: we imagine that Φ is a substitution of variables and take differentials. But then we write it up as a formal pullback so that we can maintain our smug attitude towards physicists.

It's not totally unreasonable to think of the pullback as some sort of "Jacobian for forms." Just like the Jacobian matrix of Φ linearizes Φ in order to tell us the effect on a given vector, the pullback is a linearization of Φ which tells us the effect of Φ on a k-form, except that we pull forms backwards rather than push them forward. Also like the Jacobian, the pullback plays nicely with products and derivatives. For any open subset $\mathcal{R} \subset \mathbb{R}^n$ and differential forms $\eta, \lambda \in \Omega^\bullet(\mathcal{R})$, we have

$$\boxed{\Phi^*(\eta \wedge \lambda) = (\Phi^*\eta) \wedge (\Phi^*\lambda)} \quad \text{and} \quad \boxed{\Phi^*(d\eta) = d(\Phi^*\eta).} \tag{9.2}$$

In some sense, these equations verify that, however Φ may distort differential forms, this distortion is accounted for by the pullback. Our final two examples verify (9.2). In order to verify the wedge product formula, we will use the following algebraic fact.

Exercise 9.15. Show that if $\lambda^1, \dots, \lambda^k$ are 1-forms on an open subset of $\mathbb{R}^n$, then

$$(\lambda^1 \wedge \cdots \wedge \lambda^k)(\mathbf{v}_1, \dots, \mathbf{v}_k) = \det \begin{pmatrix} \lambda^1(\mathbf{v}_1) & \cdots & \lambda^1(\mathbf{v}_k) \\ \vdots & \ddots & \vdots \\ \lambda^k(\mathbf{v}_1) & \cdots & \lambda^k(\mathbf{v}_k) \end{pmatrix},$$

for any vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$ in $\mathbb{R}^n$.

Example 9.16. Let's verify the first part of (9.2), which tells us that pullback plays nicely with wedge product. Exercise 9.12 verifies this in the case that one of the parties to our wedge product is a 0-form, and now we consider a higher-order wedge product of 1-forms. Suppose that $I = (i_1, \dots, i_k)$ is a multi-index of integers between 1 and n. Then we find that

$$\begin{aligned}(\Phi^* dx^I)(\mathbf{v}_1, \dots, \mathbf{v}_k) &= dx^I(\Phi_*\mathbf{v}_1, \dots, \Phi_*\mathbf{v}_k) = (dx^{i_1} \wedge \cdots \wedge dx^{i_k})(\Phi_*\mathbf{v}_1, \dots, \Phi_*\mathbf{v}_k) \\ &= \det \begin{pmatrix} dx^{i_1}(\Phi_*\mathbf{v}_1) & \cdots & dx^{i_1}(\mathbf{v}_k) \\ \vdots & \ddots & \vdots \\ dx^{i_k}(\Phi_*\mathbf{v}_1) & \cdots & dx^{i_k}(\mathbf{v}_k) \end{pmatrix} = \det \begin{pmatrix} (\Phi^* dx^{i_1})(\mathbf{v}_1) & \cdots & (\Phi^* dx^{i_1})(\mathbf{v}_k) \\ \vdots & \ddots & \vdots \\ (\Phi^* dx^{i_k})(\mathbf{v}_1) & \cdots & (\Phi^* dx^{i_k})(\mathbf{v}_k) \end{pmatrix} \\ &= ((\Phi^* dx^{i_1}) \wedge \cdots \wedge (\Phi^* dx^{i_k}))(\mathbf{v}_1, \dots, \mathbf{v}_k).\end{aligned}$$

In case you don't believe the third equality, we point out that the j^{th} row of the matrix whose determinant it computes is the $i_j{}^{\text{th}}$ row of the matrix whose columns are $\Phi_*\mathbf{v}_1, \dots, \Phi_*\mathbf{v}_k$. The final equality uses the result of Exercise 9.15. So we have verified that $\Phi^* dx^I = (\Phi^* dx^{i_1}) \wedge \cdots \wedge (\Phi^* dx^{i_k})$. The more general property given by (9.2) then follows from the following exercise, which is solved using properties of the wedge product.

Exercise 9.17. Using the above formula for $\Phi^* dx^I$, show that $\Phi^*(dx^I \wedge dx^J) = (\Phi^* dx^I) \wedge (\Phi^* dx^J)$, for any appropriate multi-indices I and J. Then show that $\Phi^*(\eta \wedge \lambda) = (\Phi^*\eta) \wedge (\Phi^*\lambda)$ for any forms η, λ on $\mathcal{R}$.

Exercise 9.18. Using the map $\Phi\colon \mathbb{R}^3_{r,\theta,z} \to \mathbb{R}^3_{x,y,z}$ from Example 9.14, show that

$$\Phi^*(dx \wedge dy \wedge dz) = r\, dr \wedge d\theta \wedge dz.$$

This is reminiscent of the "distortion factor" induced by Φ back in part one of the course.

Exercise 9.19. Consider the map $\Phi\colon \mathbb{R}^3_{\rho,\theta,\varphi} \to \mathbb{R}^3_{x,y,z}$ defined by

$$\Phi(\rho, \theta, \varphi) := (\rho \sin\varphi \cos\theta, \rho \sin\varphi \sin\theta, \rho \cos\varphi).$$

Make a conjecture about the pullback $\Phi^*(dx \wedge dy \wedge dz)$, and then compute this pullback.

Finally, we verify the second boxed formula in (9.2), which tells us that pullback is compatible with differentiation.

Example 9.20. Consider a k-form η on an open subset $\mathcal{R}$ of $\mathbb{R}^n$, along with a smooth function $\Phi\colon \mathcal{D} \to \mathcal{R}$, for some open subset $\mathcal{D}$ of $\mathbb{R}^m$. We want to show that $\Phi^*(d\eta) = d(\Phi^*\eta)$. Say we have a point p in $\mathcal{D}$, along

with vectors $\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}$ based at p, which we treat as constant vector fields on $\mathcal{D}$. For each $1 \leq j \leq k+1$, let's define a function $f_j : \mathcal{R} \to \mathbb{R}$ via

$$f_j(q) := \eta_q(\Phi_* \mathbf{v}_1, \ldots, \widehat{\Phi_* \mathbf{v}_j}, \ldots, \Phi_* \mathbf{v}_{k+1}).$$

Notice that

$$d(\Phi^* \eta)_p(\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}) = \sum_{j=1}^{k+1} (-1)^{j+1} D_{\mathbf{v}_j}(f_j \circ \Phi)\Big|_p.$$

On the other hand, we find that

$$\begin{aligned}(\Phi^*(d\eta))_p(\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}) &= (d\eta)_{\Phi(p)}(\Phi_* \mathbf{v}_1, \ldots, \Phi_* \mathbf{v}_k) \\ &= \sum_{j=1}^{k+1} (-1)^{j+1} D_{\Phi_* \mathbf{v}_j} f_j\Big|_{\Phi(p)} = \sum_{j=1}^{k+1} (-1)^{j+1} D_{\mathbf{v}_j}(f_j \circ \Phi)\Big|_p,\end{aligned}$$

where the last equality uses the chain rule for directional derivatives:

$$D_{\mathbf{v}}(f \circ \Phi)\Big|_p = D_{\Phi_* \mathbf{v}} f\Big|_{\Phi(p)},$$

for any smooth function $f : \mathcal{R} \to \mathbb{R}$. So indeed, $\Phi^*(d\eta) = d(\Phi^* \eta)$.

The pullback will be crucial to our definition of integration over parametrized subsets of $\mathbb{R}^n$. Loosely, it's straightforward to integrate an n-form over an open subset of $\mathbb{R}^n$: there is a natural way to treat an n-form on $\mathbb{R}^n$ as a real-valued function, and we'll just integrate this function as in part one of the course. But if we want to integrate a k-form over a k-dimensional subset $\mathcal{S}$ of $\mathbb{R}^n$ — with $k < n$ — we'll need a parametrization $\Phi : \mathcal{D} \to \mathbb{R}^n$, where $\mathcal{D}$ is an open subset of $\mathbb{R}^k$. (The existence of such a parametrization will more or less be what we mean when we say that $\mathcal{S}$ is k-dimensional.) We can then realize our integral over $\mathcal{S}$ as an integral over $\mathcal{D}$ via the pullback Φ^*.

9.3 Day 25: Integrating k-forms

Goals

By the end of today's class, we should be able to do the following.

1. Integrate **top-degree forms** over certain subsets of Euclidean space.
2. Integrate differential k-forms over **parametrized k-surfaces**.

Most mathematicians, when asked, "What's a differential form?" will answer, "It's a thing to be integrated." My own feeling is that this is a bit unfair to differential forms: they're useful not only for integral calculus, but also *differential* calculus, and allow us to express partial differential relations in coordinate-free ways. But today we overlook this objection and think about what it means to integrate a differential form.

9.3.1 Top-degree forms

We start with integration of differential forms over open subsets of $\mathbb{R}^n$. Our first observation is that integration in this setting is defined only for **top-degree forms**.

Definition. Let $\mathcal{R}$ be an open subset of $\mathbb{R}^n$, and let λ be a differential form on $\mathcal{R}$. We call λ a **top-degree form** on $\mathcal{R}$ if λ is an n-form.

Top-degree forms are particularly special because, up to multiplication by elements of $C^\infty(\mathcal{R})$, there's just one. That is, any top-degree form λ on an open subset $\mathcal{R}$ of $\mathbb{R}^n$ can be written as

$$\lambda = f \, dx^1 \wedge \cdots \wedge dx^n,$$

for some smooth function $f \in C^\infty(\mathcal{R})$. The upshot is that, given the coordinates $x_1, \ldots, x_n$, we can identify the set of top-degree forms on $\mathcal{R}$ with the set of smooth functions on $C^\infty(\mathcal{R})$. But we already know how to integrate these!

Definition. Let $\mathcal{R}$ be an compact subset of $\mathbb{R}^n$ which has nonempty interior, and whose boundary has volume zero[2], and let λ be a differential n-form on $\mathring{\mathcal{R}}$. Then the **integral of λ over $\mathcal{R}$** is given by

$$\int_{\mathcal{R}} \lambda := \int_{\mathcal{R}} f,$$

where $\lambda = f\, dx^1 \wedge \cdots \wedge dx^n$, and the integral on the right is defined as in part one of our course.

Remark. Back in part one of the course, we'd be wringing our hands over whether or not f is integrable over $\mathcal{R}$. But remember that in this last part of the course, all functions are assumed to be infinitely differentiable. In particular, f is continuous, and thus integrable. One can think about what happens if f is not assumed to be smooth, but we won't concern ourselves with this right now.

Honestly, this definition probably feels like a bit of a letdown. Why go through the pain of defining differential forms if integrating them is really just the same as integrating functions? One glib-sounding answer is that we were never really integrating functions — the forms were there the entire time. This has some truth to it, and we try to give a hint as to what this means in Section 9.3.4 below. Another reason to integrate forms rather than functions is that integration of functions is not invariant under coordinate change, while coordinate-invariance is built into the definition of differential forms.

To see what we mean, suppose we have a smooth, bijective, regular[3] map $\Phi\colon \mathcal{D} \to \mathcal{R}$ between appropriate subsets of $\mathbb{R}^n$. We can reasonably think of Φ as giving us a new coordinate system on $\mathcal{R}$ — compare Φ with the coordinate systems from Week 3 — and thus we can integrate f against the coordinates of $\mathcal{D}$. It doesn't seem far-fetched then, that we would want to the value given by the integral

$$\int_{\mathcal{D}} f \circ \Phi$$

to agree with $\int_{\mathcal{R}} f$, since Φ is just a "reparametrization" of the $\mathcal{R}$. But of course we know that this fails — we spent a couple of days earlier this quarter thinking about the change of variables theorem, which tells us how to correct the function $f \circ \Phi$ with a distortion factor. The notion of composition for differential forms is given by the pullback, and what we want to see now is that the distortion factor is built into the pullback.

Theorem 9.21: Pullback of top-degree forms

Let $\Phi\colon \mathcal{D} \to \mathcal{R}$ be a smooth map between open subsets of $\mathbb{R}^n$, and consider coordinates $u_1, \ldots, u_n$ on $\mathcal{D}$ and coordinates $x_1, \ldots, x_n$ on $\mathcal{R}$. Then

$$\Phi^*(dx^1 \wedge \cdots dx^n) = \det \Phi_*\, du^1 \wedge \cdots \wedge du^n,$$

where Φ_* is the $n \times n$ Jacobian matrix of Φ.

Proof. To ease our notational burden slightly, let's define the multi-index $I = (1, \ldots, n)$, so that

$$dx^I = dx^1 \wedge \cdots \wedge dx^n \quad \text{and} \quad du^I = du^1 \wedge \cdots \wedge du^n.$$

Now consider a collection of vectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$ in $\mathbb{R}^n$, thought of as tangent vectors on the domain side of Φ. Unwinding definitions gives

$$\begin{aligned}(\Phi^* dx^I)(\mathbf{v}_1, \ldots, \mathbf{v}_n) = dx^I(\Phi_*\mathbf{v}_1, \ldots, \Phi_*\mathbf{v}_n) &= \det \begin{pmatrix} | & & | \\ \Phi_*\mathbf{v}_1 & \cdots & \Phi_*\mathbf{v}_n \\ | & & | \end{pmatrix} \\ &= \det \left(\Phi_* \begin{pmatrix} | & & | \\ \mathbf{v}_1 & \cdots & \mathbf{v}_n \\ | & & | \end{pmatrix} \right) = \det \Phi_* \cdot \det \begin{pmatrix} | & & | \\ \mathbf{v}_1 & \cdots & \mathbf{v}_n \\ | & & | \end{pmatrix} \\ &= \det \Phi_*\, du^I(\mathbf{v}_1, \ldots, \mathbf{v}_n).\end{aligned}$$

[2]These conditions can be recalled from part one of the course.

[3]Meaning that the Jacobian determinant is nowhere zero.

Because this computation holds for any collection of vectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$, we find that $\Phi^* dx^I = \det \Phi_* \, du^I$, as desired. □

Remark. We could alternatively prove Theorem 9.21 using the fact that pullback distributes over wedge product, as well as some algebraic facts about the determinant. Thinking about that proof is probably not a bad use of your time; the approach used here was meant to emphasize that top-degree forms compute volumes, and the Jacobian determinant of Φ tells us how Φ distorts these volumes.

You should compare the result of Theorem 9.21 with your solutions to Exercises 9.18 and 9.19.

Combined with Theorem 3.5, Theorem 9.21 tells us that pullback is the correct notion of composition for differential forms. Namely, when we pull our integrals back over smooth, orientation-preserving maps, we simply pull back our forms as well.

Theorem 9.22: The change of variables theorem for forms

Let $\Phi \colon \mathcal{D} \to \mathcal{R}$ be a smooth map between compact subsets of $\mathbb{R}^n$ with nonempty interiors, and whose boundaries have volume zero. Let λ be a differential n-form on $\mathring{\mathcal{R}}$. If the Jacobian determinant of Φ is positive[4]at all points of $\mathring{\mathcal{D}}$, then

$$\int_{\mathcal{D}} \Phi^* \lambda = \int_{\mathcal{R}} \lambda.$$

Exercise 9.23. Prove Theorem 9.22.

Perhaps this still seems like a poor use of our time — we've spent so many hours developing differential forms so that we can compute a pullback instead of a Jacobian determinant? All this to move a star from the bottom of Φ to the top? The importance of this coordinate-invariance for integration of differential forms will become clearer in the next subsection, where we integrate over parametrized surfaces.

9.3.2 Parametrized surfaces

We now give a quick definition of some objects over which we can integrate differential forms.

Definition. Let $\mathcal{R}$ be an open subset of $\mathbb{R}^n$. For any integer $k \geq 0$, a **parametrized k-surface in** $\mathcal{R}$ is a smooth map

$$\Phi \colon \mathcal{D} \to \mathcal{R},$$

where $\mathcal{D}$ is a compact subset of $\mathbb{R}^k$ with nonempty interior whose boundary has volume zero. We call $\mathcal{D}$ the **parameter domain of** Φ**.**

Remark. It's important to notice that the object we're calling a parametrized surface is the map itself — not the image of the map. This may not always match our intuition, but will be an important point in writing down integrals.

Example 9.24. Say we have a parameter domain $\mathcal{D} \subset \mathbb{R}^k$ — that is, a compact subset of $\mathbb{R}^k$ with nonempty interior whose boundary has volume zero. Then the inclusion map $\iota \colon \mathcal{D} \to \mathbb{R}^k$, defined by

$$\iota(x) := x,$$

is a parametrized k-surface in $\mathbb{R}^k$.

Example 9.25. A parametrized 0-surface in $\mathcal{R} \subset \mathbb{R}^n$ is simply the inclusion of a point. By convention, $\mathbb{R}^0$ is a set containing a single point — we often write $\mathbb{R}^0 = \{*\}$ — and thus there is only one parameter domain in $\mathbb{R}^0$: all of $\mathbb{R}^0$. So a parametrized 0-surface is a smooth map

$$\Phi \colon \mathbb{R}^0 \to \mathcal{R}$$

which is to say any map $\Phi \colon \{*\} \to \mathcal{R}$. (The smooth requirement is treated as vacuous here.)

[4]When $\det \Phi_* > 0$, we say that Φ is **orientation-preserving**.

It's important to note that we do not require our parametrized k-surfaces to be simple or regular in any way. Because we focus on the data of the map itself — rather than the image of the map — we will be able to define integration without caring about whether or not the k-surface self-intersects or fails at some points to be regular.

9.3.3 Integration over parametrized surfaces

We've now set the stage for integration of differential k-forms over parametrized k-surfaces.

A parametrized k-surface is a map whose image lives in an open subset $\mathcal{R}$ of $\mathbb{R}^n$, and we imagine[5] this image to be k-dimensional. So it makes sense that we would want to integrate elements of $\Omega^k(\mathcal{R})$ over our k-surface. Theorem 9.22 tells us that it should be safe enough to simply pull a k-form on $\mathcal{R}$ back to a k-form on $\mathcal{D}$.

Definition. With $k \geq 0$ fixed, let $\Phi\colon \mathcal{D} \to \mathcal{R}$ be a parametrized k-surface in $\mathbb{R}^n$, for some open subset $\mathcal{R}$ of $\mathbb{R}^n$. For any differential k-form λ on $\mathcal{R}$, we define the **integral of λ over Φ** to be

$$\int_\Phi \lambda := \int_{\mathcal{D}} \Phi^* \lambda.$$

Exercise 9.26. Convince yourself that integrating a 0-form over a parametrized 0-surface simply amounts to evaluating a function at a point.

Integration over parametrized k-surfaces is meant to generalize the line integrals and surface integrals that we computed in part two of the course. Indeed, while we pretended to define line integrals by doing some sort of partitioning-the-curve nonsense, we defined the surface integral of a scalar function using parametrizations of the surface. (Really, we did! Go back and look at Equation 6.2 on page 74.) But our fear of computing areas of parallelograms in dimensions greater than three restricted us[6] to surfaces in $\mathbb{R}^3$, and this parametrization-focused approach forced us to include an ugly distortion factor in our definition. Our forms-based approach builds the distortion factors into our definition of composition; composition is encapsulated by pullback, and distortion factors are baked into the pullback.

At the risk of repeating ourselves, let's think about what happens if we try to define integration of functions over parametrized k-surfaces in the way that seems most natural. Given a parametrized k-surface $\Phi\colon \mathcal{D} \to \mathcal{R}$ and a smooth function $f : \mathcal{R} \to \mathbb{R}$, it seems perfectly natural to expect the integral of f over Φ to be given by

$$\int_{\mathcal{D}} f \circ \Phi,$$

since $f \circ \Phi\colon \mathcal{D} \to \mathbb{R}$ is the pullback of f to $\mathcal{D}$. The problem is that this integral is very sensitive to reparametrization. Say we have a smooth, bijective map $\Psi\colon \mathcal{D}' \to \mathcal{D}$, for some other parameter domain $\mathcal{D}' \subset \mathbb{R}^k$. We think of Ψ as a reparametrization; not only does $\Phi \circ \Psi$ parametrize the same k-surface as does Φ, it preserves any simplicity, regularity, or lack thereof enjoyed by Φ. So we should expect the integral of f over $\Phi \circ \Psi$ to be computed by

$$\int_{\mathcal{D}'} f \circ (\Phi \circ \Psi).$$

But of course this can't be so. We know that

$$\int_{\mathcal{D}} f \circ \Phi = \int_{\mathcal{D}'} (f \circ \Phi \circ \Psi) |\det \Psi_*|,$$

so we are essentially asking $\det \Psi_*$ to be everywhere equal to ± 1. If we want to integrate functions over parametrized k-surfaces, we are forced to incorporate the distortion factors of our parametrizations into the definition. There is a way to do this without differential forms — the Gramian determinant of Exercise 8.26 makes an appearance — but the machinery of differential forms gives us an elegant way to package all of these distortion factors into one algebraic gadget.

[5] There are some problems with this interpretation of dimension, but we'll not address these here. The problems are tied up with the fact that our parametrized k-surface need not be simple or regular.

[6] I think Fred may have discussed surfaces in higher dimensions in discussion session, though.

Example 9.27. Consider the 2-form $\nu \in \Omega^2(\mathbb{R}^3)$ defined by

$$\nu = x\,dy \wedge dz - y\,dx \wedge dz + z\,dx \wedge dy,$$

and consider the parametrized 2-surface $\Phi\colon [0,\pi]_\varphi \times [0,2\pi]_\theta \to \mathbb{R}^3$ defined by

$$\Phi(\varphi,\theta) := (\sin\varphi\cos\theta, \sin\varphi\sin\theta, \cos\varphi).$$

Let's compute $\int_\Phi \nu$. We start by computing the pullback of ν under Φ. We have

$$\begin{aligned}
\Phi^* dx &= \frac{\partial}{\partial\varphi}(\sin\varphi\cos\theta)\,d\varphi + \frac{\partial}{\partial\theta}(\sin\varphi\cos\theta)\,d\theta = \cos\varphi\cos\theta\,d\varphi - \sin\varphi\sin\theta\,d\theta \\
\Phi^* dy &= \frac{\partial}{\partial\varphi}(\sin\varphi\sin\theta)\,d\varphi + \frac{\partial}{\partial\theta}(\sin\varphi\sin\theta)\,d\theta = \cos\varphi\sin\theta\,d\varphi + \sin\varphi\cos\theta\,d\theta \\
\Phi^* dz &= \frac{\partial}{\partial\varphi}(\cos\varphi)\,d\varphi + \frac{\partial}{\partial\theta}(\cos\varphi)\,d\theta = -\sin\varphi\,d\varphi.
\end{aligned}$$

From this we can compute

$$\begin{aligned}
\Phi^*(dy\wedge dz) &= \sin^2\varphi\cos\theta\,d\varphi\wedge d\theta \\
\Phi^*(dx\wedge dz) &= -\sin^2\varphi\sin\theta\,d\varphi\wedge d\theta \\
\Phi^*(dx\wedge dy) &= \cos\varphi\sin\varphi\,d\varphi\wedge d\theta,
\end{aligned}$$

and at last we see that

$$\begin{aligned}
\Phi^*\nu &= (x\circ\Phi)\Phi^*(dy\wedge dz) - (y\circ\Phi)\Phi^*(dx\wedge dz) + (z\circ\Phi)\Phi^*(dx\wedge dy) \\
&= (\sin\varphi\cos\theta)(\sin^2\varphi\cos\theta)\,d\varphi\wedge d\theta + (\sin\varphi\sin\theta)(\sin^2\varphi\sin\theta)\,d\varphi\wedge d\theta \\
&\qquad + (\cos\varphi)(\cos\varphi\sin\varphi)\,d\varphi\wedge d\theta \\
&= (\sin^3\varphi\cos^2\theta + \sin^3\varphi\sin^2\theta + \sin\varphi\cos^2\varphi)\,d\varphi\wedge d\theta \\
&= \sin\varphi\,d\varphi\wedge d\theta.
\end{aligned}$$

So we find that

$$\int_\Phi \nu = \int_{[0,\pi]\times[0,2\pi]} \Phi^*\nu = \int_{[0,\pi]\times[0,2\pi]} \sin\varphi\,d\varphi\wedge d\theta.$$

Now $d\varphi\wedge d\theta$ is the standard top-degree form on $[0,\pi]_\varphi\times[0,2\pi]_\theta$, so this last integral simply asks us to compute the integral (in the sense of part one of the course) of $\sin\varphi$. That is,

$$\int_\Phi \nu = \int_{[0,\pi]\times[0,2\pi]} \sin\varphi\,d\varphi\wedge d\theta = \int_{[0,\pi]\times[0,2\pi]} \sin\varphi = \cdots = \boxed{4\pi}.$$

Exercise 9.28. (**Challenge**) Suppose that $\Phi\colon \mathcal{D}\to\mathbb{R}^3$ is a simple, regular, parametrized 2-surface in $\mathcal{R}^3$, and let $\mathcal{S}$ denote the image of Φ. With coordinates (u,v) on $\mathcal{D}$, let

$$\mathbf{n}(u,v) := \frac{\Phi_u\times\Phi_v}{\|\Phi_u\times\Phi_v\|}$$

denote the oriented unit normal vector to $\mathcal{S}$ determined by Φ. If we pretend that $\mathbf{n}(u,v) = \langle n_1, n_2, n_3\rangle$ is defined on all of $\mathbb{R}^3$ (rather than just on $\mathcal{S}$), we can define a 2-form

$$\nu_{\mathcal{S}} := n_1\,dy\wedge dz - n_2\,dx\wedge dz + n_3\,dx\wedge dy$$

on $\mathbb{R}^3$. Under this pretense, prove that

$$\iint_{\mathcal{S}} f\,dS = \int_\Phi f\;\nu_{\mathcal{S}}$$

for any smooth function $f:\mathcal{S}\to\mathbb{R}$, where the integral on the left is defined as in week 6. In light of this exercise, explain why our answer of 4π makes sense in Example 9.27.

Exercise 9.29. (**Challenge**) Exercise 9.28 tells us how to interpret scalar surface integrals in terms of differential forms, for surfaces in $\mathbb{R}^3$. Under the same assumptions as in Exercise 9.28, let $\mathbf{F} = \langle F_1, F_2, F_3\rangle$ be a smooth vector field on $\mathbb{R}^3$, and define

$$\omega_{\mathbf{F}} := F_1\,dy\wedge dz - F_2\,dx\wedge dz + F_3\,dx\wedge dy.$$

Prove that

$$\iint_{\mathcal{S}} \mathbf{F}\cdot d\mathbf{S} = \int_\Phi \omega_{\mathbf{F}},$$

where, once again, the integral on the left is defined as in week 6.

9.3.4 Why forms?

In this subsection, we'll briefly try to convince ourselves that differential forms have been lurking in the background of our integrals all along. There isn't really any new material here, and certainly nothing for which you're directly responsible in this class. But hopefully this informal discussion helps to motivate the integration we introduced this week.

The very first integral you ever defined was probably the definite integral of a continuous function on a closed interval, say $f : [a, b] \to \mathbb{R}$. The purpose of the integral was to measure the signed area between the graph $y = f(x)$ and the x-axis, bounded by the lines $x = a$ and $x = b$. Of course this area will depend on the (varying) height of $y = f(x)$, as well as the length of the interval $[a, b]$. In particular, the definite integral $\int_a^b f(x)\, dx$ depends on a choice of function f, and also on some notion of length for the interval $[a, b]$.

On the other hand, the area we compute should not be affected by the names we give to the points in the interval $[a, b]$. That is, if we use some nonstandard coordinate system on $[a, b]$, the value of the integral shouldn't change, provided we're still using the same ruler — i.e., notion of length — from before.

Let's give some names to the players here. The integral over which we're trying to integrate is $I := [a, b]$, and let's say we have a coordinate system

$$u : [c, d] \to I.$$

Here we're thinking of I as an abstract geometric object over which we want to integrate; this object has a preferred notion of length for (tangent) vectors, but has no *a priori* preferred coordinate system. We're going to use the coordinate system u, and the value we compute for the integral shouldn't depend on u.

The ruler on I is not particularly interesting: it's just the standard 1-form $dx : \mathbb{R} \to \mathbb{R}$. Tangent vectors to points in I live in $\mathbb{R}$, which has the standard basis vector ∂_x, and the length form — which we'll call λ — on I is given by

$$\lambda(\mathbf{v}) := dx(\mathbf{v}) = \mathbf{v} \cdot \partial_x.$$

Of course this is overkill if we're just trying to reproduce the integrals you computed in your first calculus class, but the point is that our definition of the integral of f over I depends on a preferred notion of signed length on I.

As we're well aware, we approximate the integral of f over I by partitioning I into many subintervals, approximating the value of f over each of these subintervals, and then summing the resulting areas. One complication here is that we're treating I as some abstract geometric object; so the only reasonable way to subdivide I is using a coordinate system. For instance, we have the coordinate system $u : [c, d] \to I$, which has parameter domain $[c, d]$. We can choose a partition

$$\mathcal{P} = \{c = t_0 < t_1 < \cdots < t_N = d\}$$

of the parameter domain, for some large integer $N \gg 1$, and this leads to a subdivision of I defined by

$$I_j := u([t_{j-1}, t_j]), \quad \text{for } 1 \le j \le N.$$

If $\mathcal{P}$ is a very fine partition, then we should have

$$\int_I f\, dx \approx \sum_{j=1}^{N} f(u(t_{j-1}))\, \text{length}(I_j).$$

Here we're approximating the integral with a left-hand Riemann sum, and pretending that we can't easily compute the length of the subinterval of I_j.

The remaining question is how we compute the lengths of the subintervals I_j. Given what we know about I, this is easy to do, but we're trying to treat I as an abstract geometric object where we know how to compute lengths of vectors (using λ), but not of subsets. Our plan, then, is to approximate the subintervals I_j with vectors tangent to I. In the parameter domain of u we can consider the vectors

$$\mathbf{v}_j := (t_j - t_{j-1})\, \partial_t, \quad \text{for } 1 \le j \le N.$$

We treat $\mathbf{v}_j$ as based at $t_{j-1} \in [c, d]$, so that $\mathbf{v}_j$ spans the subinterval $[t_{j-1}, t_j]$ of $[c, d]$. Our plan, then, is to push $\mathbf{v}_j$ forward from $[c, d]$ to I. Once we have a tangent vector to I, we can use λ to measure its length.

Since $\mathbf{v}_j$ is a vector based at t_{j-1}, there's really just one way to push it forward via u: we consider the vector $(u_*)_{t_{j-1}}\mathbf{v}_j$, where $(u_*)_{t_{j-1}}$ denotes the Jacobian matrix u_* of u, evaluated at t_{j-1}. So, altogether, we find that

$$\text{length}(I_j) \approx \lambda((u_*)_{t_{j-1}}\mathbf{v}_j).$$

Of course the expression on the right is just the pullback form $u^*\lambda$, evaluated at the point t_{j-1} and on the vector $\mathbf{v}_j$. So

$$\int_I f\,dx \approx \sum_{j=1}^{N} f(u(t_{j-1}))(u^*\lambda)_{t_{j-1}}(\mathbf{v}_j) = \sum_{j=1}^{N} f(u(t_{j-1}))(u^*\lambda)_{t_{j-1}}(\partial_t)(t_j - t_{j-1}).$$

At this point we recognize the final expression as a left-hand Riemann sum for the function $g : [c,d] \to \mathbb{R}$ defined by

$$g(t) := f(u(t))(u^*\lambda)_t(\partial_t).$$

In particular, this is a left-hand Riemann sum using the standard coordinate system on $[c,d]$. So this means that

$$\int_I f\,dx = \int_{[c,d]} g\,dt = \int_{[c,d]} (f\circ u)\,u^*\lambda = \int_{[c,d]} u^*(f\,dx).$$

In the language of differential forms, we see that we weren't integrating the function f, with dx just riding along as notation. Instead, we were integrating the 1-form $f\,dx$. The value of this integral is independent of our coordinate system, provided we transform the 1-form $f\,dx$ appropriately.

Said another way, the 1-form dx gives us our ruler — a way to measure lengths — and this ruler is unchanged under coordinate transformations (though the way we write the ruler will change). We can't compute areas under curves without some notion of length, and this 1-form is precisely what gives us our lengths.

This probably still feels pretty fishy. The best way to get your head around the ideas that appear here is to consider even more cases.

Exercise 9.30. Say we have an oriented curve $\mathcal{C} \subset \mathbb{R}^n$ which admits a simple, regular, positively-oriented parametrization. Under these assumptions, it's possible to write down a 1-form λ on $\mathbb{R}^n$ so that $\lambda_p(\mathbf{v})$ gives the signed magnitude of $\mathbf{v}$, for every point $p \in \mathcal{C}$ and every vector $\mathbf{v}$ tangent to $\mathcal{C}$ at p. Given a simple, regular parametrization $\mathbf{r}\colon [a,b] \to \mathbb{R}^2$ of $\mathcal{C}$, run through the above discussion in order to define the integral over $\mathcal{C}$ of a continuous function $f : \mathcal{C} \to \mathbb{R}$. Convince yourself that all we've done is integrate the pullback of the 1-form λ over $[a,b]$.

Just as integration in one dimension requires some notion of (signed) length, integration in two dimensions requires a notion of (signed) area.

Exercise 9.31. Consider a surface $\mathcal{S} \subset \mathbb{R}^n$, assuming that at every point $p \in \mathcal{S}$ we can make sense of the tangent space to $\mathcal{S}$ at p, and that this tangent space is two-dimensional. Suppose we have a 2-form ω defined on $\mathbb{R}^n$ with the property that $\omega_p(\mathbf{v},\mathbf{w})$ measures the signed area of the parallelogram spanned by $\mathbf{v}$ and $\mathbf{w}$, for every $p \in \mathcal{S}$ and vectors $\mathbf{v},\mathbf{w}$ tangent to $\mathcal{S}$ at p. (Don't bother trying to write down an expression for λ — just take it as something that exists.) Assuming we have a bijective parametrization

$$\Phi\colon \mathcal{D} \to \mathbb{R}^n$$

of $\mathcal{S}$, work through the above argument to convince yourself that the surface integral $\int_{\mathcal{S}} f\,dS$ is simply the integral $\int_{\mathcal{D}} \Phi^*(f\,\omega)$ of the pullback of $f\,\omega$ over $\mathcal{D}$, for any continuous function $f : \mathcal{S} \to \mathbb{R}$.

Note. Our mantra throughout this section was that integration over a closed interval requires both a function and a notion of length on the interval, and we then found that the length was encoded by a differential form. This undersells differential forms a bit. It's true that we need a notion of length for our domain when writing down the definite integral of a function on a closed interval, and we need a notion of area for our domain when integrating functions of two variables. It's also true that the framework of differential forms was developed to give us a single, cohesive algebraic framework for these various notions of integration. But this framework is sufficiently robust that we can write down integrals which have nothing to do with length or area. In a setting where we can't talk about the lengths of vectors, it will be impossible to speak meaningfully about the integral of a function, but we can still integrate differential forms. Maybe right now it's not clear why such a notion of integration would be useful, but many notions of *action* in physics are defined by integrals which do not necessarily care about underlying lengths.

9.4 Interlude: The algebra (and calculus) of differential forms on $\mathbb{R}^n$

Our approach to the algebra of differential forms on (open subsets of) $\mathbb{R}^n$ was to write everything down in terms of standard k-forms dx^I. This made our definitions feel somewhat less abstract, but may have obfuscated some of the important algebraic properties we use when actually manipulating forms.

Our definition of the exterior derivative is also not particularly useful for actual computation. This is similar to the single-variable calculus case, where we write down a definition of the derivative that computes a limit, but then quickly learn rules that allow us to compute a large class of derivatives without directly thinking about the definition.

This bonus section summarizes many properties of differential forms and the various operations on them that we've defined. You should be able to verify any of these properties using the definitions given above[7]; indeed, several appeared as exercises. Also, the properties are not independent of one another: some are special cases of other properties on the list. We make no claim that the list is exhaustive.

Throughout, $\mathcal{R}$ is an open subset of $\mathbb{R}^n$, p is a point in $\mathcal{R}$, λ and η are k- and ℓ-forms on $\mathcal{R}$, respectively, and $\mathbf{v}_i$ is either a vector based at p or a vector field defined on $\mathcal{R}$, for any subscript i. We also have a smooth map $\Phi\colon \mathcal{D} \to \mathcal{R}$, where $\mathcal{D}$ is an open subset of $\mathbb{R}^m$.

- (forms are multi-linear) For any auxiliary vector field $\mathbf{w}$ on $\mathcal{R}$,

  $$\lambda(\mathbf{v}_1, \dots, \mathbf{v}_i, \dots, \mathbf{v}_k) + \lambda(\mathbf{v}_1, \dots, \mathbf{w}, \dots, \mathbf{v}_k) = \lambda(\mathbf{v}_1, \dots, \mathbf{v}_i + \mathbf{w}, \dots, \mathbf{v}_k),$$

  where $\mathbf{w}$ is in the i^{th} position.
- (forms are alternating) If the i^{th} and j^{th} vectors in the expression $\lambda(\mathbf{v}_1, \dots, \mathbf{v}_k)$ are swapped, the output is scaled by -1.
- (forms are alternating, again) If the vectors $\mathbf{v}_1, \dots, \mathbf{v}_k$ are linearly dependent, then $\lambda(\mathbf{v}_1, \dots, \mathbf{v}_k) = 0$. In particular, if $\mathbf{v}_i = \mathbf{v}_j$ for some $i \neq j$, then $\lambda(\mathbf{v}_1, \dots, \mathbf{v}_k) = 0$.
- (dimension property) If $k > n$, then $\lambda = 0$.
- (wedges of standard forms) For any multi-index $I = (i_1, \dots, i_k)$ of integers between 1 and n,

  $$dx^I = dx^{i_1} \wedge \cdots \wedge dx^{i_k}.$$

- (skew-commutativity of wedge product) $\eta \wedge \lambda = (-1)^{k\ell} \lambda \wedge \eta$
- (associativity of wedge product) For any $\lambda, \eta, \theta \in \Omega^\bullet(\mathcal{R})$,

  $$(\lambda \wedge \eta) \wedge \theta = \lambda \wedge (\eta \wedge \theta).$$

- (homogeneity of wedge product) If c is a scalar, then

  $$(c\,\lambda) \wedge \eta = c\,(\lambda \wedge \eta) = \lambda \wedge (c\,\eta).$$

- (distributivity of wedge product) If η and θ have the same degree, then

  $$\lambda \wedge (\eta + \theta) = \lambda \wedge \eta + \lambda \wedge \theta.$$

- (1-forms square to zero) If λ is a 1-form, then $\lambda \wedge \lambda = 0$.
- (wedge products of 1-forms) If λ and η are 1-forms, then

  $$(\lambda \wedge \eta)(\mathbf{v}_1, \mathbf{v}_2) = \lambda(\mathbf{v}_1)\eta(\mathbf{v}_2) - \lambda(\mathbf{v}_2)\eta(\mathbf{v}_1).$$

- (derivative in terms of basis) If we write $\lambda = \sum_I f_I\, dx^I$, then

  $$d\lambda = \sum_I df_I \wedge dx^I = \sum_I \sum_{j=1}^{n} \frac{\partial f}{\partial x_j}\, dx^j \wedge dx^I.$$

  In particular, $d\lambda$ is a $(k+1)$-form.

[7] Some of the properties are actually built into the definitions.

- (d is linear) If c is any scalar and η and θ have the same degree, then

$$d(c\,\lambda) = c\,d\lambda \quad \text{and} \quad d(\eta+\theta) = d\eta + d\theta.$$

- (product rule for 0-forms) For any smooth function $f \in \Omega^0(\mathcal{R})$, $d(f\,\lambda) = df \wedge \lambda + f\,d\lambda$. In particular, $d(f\,dx^I) = df \wedge dx^I$.
- (Leibniz rule) $d(\lambda \wedge \eta) = d\lambda \wedge \eta + (-1)^k \lambda \wedge d\eta$
- (nilpotence of d) $d(d\lambda) = 0$
- (pullback on 0-forms) For any $f \in \Omega^0(\mathcal{R})$, $\Phi^* f = f \circ \Phi$.
- (pullback is linear)If c is any scalar and η and θ have the same degree, then

$$\Phi^*(c\,\lambda) = c\,\Phi^*\lambda \quad \text{and} \quad \Phi^*(\eta+\theta) = \Phi^*\eta + \Phi^*\theta.$$

- (pullback distributes over wedge products) $\Phi^*(\lambda \wedge \eta) = (\Phi^*\lambda) \wedge (\Phi^*\eta)$
- (pullback commutes with derivatives) $d(\Phi^*\lambda) = \Phi^*(d\lambda)$
- (pullback is contravariant) If $\Psi\colon \mathcal{D}' \to \mathcal{D}$ is a smooth map with basis $\mathcal{D}'$ and open subset of some Euclidean space, then $(\Phi \circ \Psi)^*\lambda = \Psi^*\Phi^*\lambda$.
- (pullback on standard 1-forms) If $\mathbb{R}^n$ has coordinates $x_1, \ldots, x_n$ and $\mathbb{R}^m$ has coordinates $u_1, \ldots, u_m$, then

$$\Phi^* dx^i = \sum_{j=1}^{m} \frac{\partial \Phi_i}{\partial u_j}\, du^j,$$

  where $\Phi = (\Phi_1, \ldots, \Phi_n)$.
- (pullback of top-degree forms) If $m = n$, then $\Phi^* dx^I = \det \Phi_*\, du^I$, where we use coordinates $u_1, \ldots, u_n$ on $\mathcal{D}$ and $x_1, \ldots, x_n$ on $\mathcal{R}$, and where $I = (1, \ldots, n)$.

Note: This list may grow/be rearranged over the next several days.

10 Week 10

This week our course reaches its crescendo: the proof of the general version of Stokes' theorem. We start by introducing **singular k-cubes**, special type of parametrized k-surface, and we learn how to take boundaries of k-cubes to produce **singular k-chains**. Then we give a proof of Stokes' theorem for k-chains. Hopefully we'll have time at the end of the week to discuss some interesting consequences of Stokes' theorem, as well as further topics you can pursue.

10.1 Day 26: Cubes and chains

Goals

By the end of today's class, we should be able to do the following.

1. Define **singular k-cubes** and **singular k-chains**.
2. Compute the **pushforward** of a singular k-chain.

Today we introduce the special kind of parametrized k-surface for which we will prove Stokes' theorem. The primary reason that we restrict our attention to the special k-surfaces is because of the (relative) ease with which we can discuss their boundaries.

10.1.1 Cubes and chains

The objects of interest to us will be **singular k-cubes**, which are simply parametrized k-surfaces whose parameter domain is the Cartesian k-cube.

Definition. For any integer $k \geq 1$, the **Cartesian k-cube** is the set $[0,1]^k$ obtained by taking the Cartesian product of $[0,1]$ with itself k times. That is,

$$[0,1]^k = \{(u_1, \ldots, u_k) : u_i \in [0,1]\}.$$

The **Cartesian 0-cube**, denoted $[0,1]^0$, is a set with exactly one element.

Remark. We'll often write $[0,1]^0 = \{*\}$, so that the unique element of $[0,1]^0$ is called $*$, but the name of this element doesn't really matter.

Definition. For any integer $k \geq 0$, the **standard k-cube** is the map $I^k : [0,1]^k \to \mathbb{R}^k$ defined by $I^k(u) := u$. For any open subset $\mathcal{R}$ of $\mathbb{R}^n$, a **singular k-cube** (or **k-cube**) **in** $\mathcal{R}$ is a smooth map $c : [0,1]^k \to \mathcal{R}$.

Remark. Often, our trademark laziness will kick in and we'll refer to singular k-cubes simply as *k-cubes*.

The adjective *singular* is present because we don't require any regularity from a singular k-cube[1]. For instance, the image in $\mathcal{R}$ of a singular k-cube could be a single point, which we typically think of as 0-dimensional rather than k-dimensional. At first this sounds like a problem, but in time we'll see that allowing degeneracies such as this ensures that the set of all singular k-cubes in $\mathcal{R}$ has nicer properties than would some more restrictive set which expects its cubes to be regular.

Exercise 10.1. Characterize all the singular 0-cubes in a subset $\mathcal{R}$ of $\mathbb{R}^n$.

Example 10.2. Consider the map $c_0 : [0,1]^2 \to \mathbb{R}^3$ defined by

$$c_0(u_1, u_2) := (\sin(\pi u_1)\cos(2\pi u_2), \sin(\pi u_1)\sin(2\pi u_2), \cos(\pi u_1)).$$

(We will typically use coordinates $u_1, \ldots, u_k$ on the Cartesian k-cube $[0,1]^k$.) Since every function in sight is smooth, so is the map c_0, and thus c_0 is a singular 2-cube. It should be clear that the image of c_0 is the unit

[1] The vocabulary kind of makes it sound like a singular k-cube is prohibited from being regular, but that's not the case. With singular k-cubes, regularity is *allowed* to fail, but not required to do so.

sphere S^2, but we think of c_0 as an object in its own right. Notice that we could also consider $c_1 : [0,1]^2 \to \mathbb{R}^3$, defined by

$$c_1(u_1, u_2) := (\sin(2\pi u_1)\cos(2\pi u_2), \sin(2\pi u_1)\sin(2\pi u_2), \cos(2\pi u_1)).$$

This will trace out the unit sphere S^2 twice — a generic point in S^2 has two preimages. But we don't think of c_1 as two copies of c_0; instead, c_0 and c_1 are independent objects. Finally, we can consider a map $\tilde{c} : [0,1]^k \to \mathbb{R}^3$, for any $k \geq 2$, defined by

$$\tilde{c}(u_1, \ldots, u_k) := (\sin(\pi u_1)\cos(2\pi u_2), \sin(\pi u_1)\sin(2\pi u_2), \cos(\pi u_1)).$$

Again, the image of $\tilde{c}$ is S^2, but this time $\tilde{c}$ is a singular k-cube, with k not necessarily equal to 2. In general, singular k-cubes can be expected to trace out subsets which have dimension at most k (but possibly smaller).

We said that we prefer singular k-cubes to more general parametrized k-surfaces because we understand their boundaries. But the boundary of a singular k-cube is not itself a singular k-cube. We solve this problem by introducing **singular k-chains**.

Definition. For any integer $k \geq 0$ and any open subset $\mathcal{R}$ of $\mathbb{R}^n$, a **singular k-chain** (or just **k-chain**) **in** $\mathcal{R}$ is a finite integer combination of singular k-cubes. That is, a singular k-chain σ in $\mathcal{R}$ has the form

$$\sigma = \sum_{j=1}^{N} n_j c_j,$$

for some $N \geq 0$, where each n_j is an integer, and each $c_j : [0,1]^k \to \mathcal{R}$ is a singular k-cube in $\mathcal{R}$. If $N = 0$ (so that our sum is empty), we have the **zero chain**. The collection of all singular k-chains on $\mathcal{R}$ is denoted $\mathcal{C}_k(\mathcal{R})$.

Remark. Though it is a linear combination of smooth maps, we do not think of a singular k-chain as a map. Instead, we imagine the singular k-cubes c_j to be the fundamental geometric objects, and a singular k-chain is then a collection of k-cubes, counted with multiplicity (the multiplicity is given by n_i).

Exercise 10.3. Consider an annulus

$$A = \{(x, y) : r^2 < x^2 + y^2 < R^2\} \subset \mathbb{R}^2,$$

for some $0 < r < R$. Express A as the image of a singular 2-cube c, but note that ∂A cannot be expressed as such an image. Write down a singular 1-chain corresponding to ∂A. We will soon be able to write down ∂c, a singular 1-chain. Keep in mind that ∂c is not the same as ∂A — one is a linear combination of smooth maps, while the other is a subset of $\mathbb{R}^2$.

Remark. Intuitively, the collection $\mathcal{C}_k(\mathcal{R})$ behaves a lot like a vector space, where the "basis elements" are the singular k-cubes in $\mathcal{R}$. Of course this means that we have infinitely many basis elements, but this isn't a problem on its own. The real hurdle to being a vector space is the fact that scalar multiplication is limited to integer values; this prevents us from doing some familiar scaling operations.

10.1.2 The pushforward

We saw above that a smooth map $\Phi : \mathcal{D} \to \mathcal{R}$ between open subsets of Euclidean spaces leads to a pullback map

$$\Phi^* : \Omega^k(\mathcal{R}) \to \Omega^k(\mathcal{D}),$$

for each $k \geq 0$. The collections of singular k-chains on $\mathcal{D}$ and $\mathcal{R}$ work similarly to their respective collections of k-forms, and instead of a pullback, Φ induces a *pushforward*.

Definition. Consider any smooth map $\Phi : \mathcal{D} \to \mathcal{R}$ between open subsets $\mathcal{D} \subset \mathbb{R}^m$ and $\mathcal{R} \subset \mathbb{R}^n$. The **pushforward map**

$$\Phi_* : \mathcal{C}_k(\mathcal{D}) \to \mathcal{C}_k(\mathcal{R})$$

is defined on singular k-cubes by

$$\Phi_* c := \Phi \circ c$$

and extended linearly to singular k-chains.

Remark. Notice that we have now overloaded the notation Φ_*. Previously this denoted the Jacobian matrix of Φ, and now it denotes the pushforward. This is not a coincidence, and maybe we will explain the relationship between these notions soon. In the meantime, our use of the notation Φ_* should be clear from context.

Exercise 10.4. Realize an arbitrary singular k-cube in $\mathcal{R}$ as a pushforward of the standard k-cube.

Recall that the pullback is *contravariant*, meaning that $(\Psi \circ \Phi)^* = \Phi^* \circ \Psi^*$. Pushforward, on the other hand, is *covariant*.

Proposition 10.5. *Suppose we have smooth maps $\Phi \colon \mathcal{D} \to \mathcal{R}$ and $\Psi \colon \mathcal{R} \to \mathcal{W}$ between open subsets of Euclidean spaces. Then $(\Psi \circ \Phi)_* = \Psi_* \circ \Phi_*$.*

Proof. If $c \colon [0,1]^k \to \mathcal{D}$ is a singular k-cube, then

$$(\Psi \circ \Phi)_* c = (\Psi \circ \Phi) \circ c = \Psi \circ (\Phi \circ c) = \Psi_*(\Phi \circ c) = \Psi_*(\Phi_* c) = (\Psi_* \circ \Phi_*)c.$$

Because pushforward is defined to be linear, it is now easy to check that $(\Psi \circ \Phi)_* \sigma = (\Psi_* \circ \Phi_*)\sigma$, for any $\sigma \in \mathcal{C}(\mathcal{D})$. So indeed, $(\Psi \circ \Phi)_* = \Psi_* \circ \Phi_*$. □

Hopefully it's clear that pushforward does for chains what pullback does for forms. By using both operations, we can express the change of variables theorem for singular k-chains. We first define integration over singular k-chains in the obvious manner.

Definition. Let $\mathcal{R}$ be an open subset of $\mathbb{R}^n$, and let σ be a singular k-chain in $\mathcal{R}$. For any $k \geq 0$ and any differential k-form on $\mathcal{R}$, the **integral of λ over σ** is given by

$$\int_\sigma \lambda := \sum_{j=1}^N n_j \int_{c_j} \lambda,$$

where $\sigma = \sum_{j=1}^N n_j c_j$.

It can be helpful to think of integration as a bilinear pairing between $\mathcal{C}_k(\mathcal{R})$ and $\Omega^k(\mathcal{R})$. We have $\langle -, - \rangle \colon \mathcal{C}_k(\mathcal{R}) \times \Omega^k(\mathcal{R}) \to \mathbb{R}$ defined by

$$\langle \sigma, \lambda \rangle := \int_\sigma \lambda.$$

Properly calling this a bilinear pairing is a little tricky, since $\mathcal{C}_k(\mathcal{R})$ only admits scalar multiplication by integers (rather than all real numbers), so we'll just use this pairing informally.

Theorem 10.6: The change of variables theorem for chains

Let $\Phi \colon \mathcal{D} \to \mathcal{R}$ be a smooth map between open subsets of Euclidean spaces, let $\sigma \in \mathcal{C}_k(\mathcal{D})$ be a singular k-chain in $\mathcal{D}$, and let $\lambda \in \Omega^k(\mathcal{R})$ be a differential k-form on $\mathcal{R}$. Then

$$\int_\sigma \Phi^* \lambda = \int_{\Phi_* \sigma} \lambda.$$

Proof. First, suppose we have a singular k-cube $c \colon [0,1]^k \to \mathcal{D}$ in $\mathcal{D}$. Then the definition of integration over k-surfaces, along with the contravariance of pullback, tells us that

$$\int_{\Phi_* c} \lambda = \int_{\Phi \circ c} \lambda = \int_{[0,1]^k} (\Phi \circ c)^* \lambda = \int_{[0,1]^k} c^* \Phi^* \lambda = \int_c \Phi^* \lambda.$$

Now if $\sigma = \sum_{j=1}^N n_j c_j$, then

$$\int_\sigma \Phi^* \lambda = \sum_{j=1}^N n_j \int_{c_j} \Phi^* \lambda = \sum_{j=1}^N n_j \int_{\Phi_* c_j} \lambda = \int_{\Phi_* \sigma} \lambda,$$

as desired. □

In terms of the integral pairing above, the change of variables theorem tells us that

$$\langle \Phi_* \sigma, \lambda \rangle = \langle \sigma, \Phi^* \lambda \rangle,$$

for any singular k-chain σ on $\mathcal{D}$ and any differential k-form λ on $\mathcal{R}$. For this reason, we like to think of Φ_* and Φ^* as **adjoint** to one another.

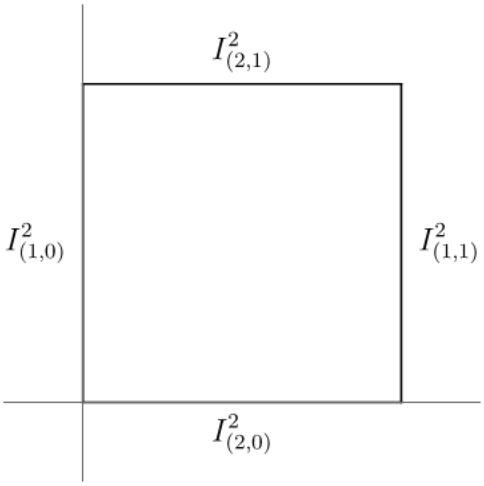

Figure 10.1: The four faces of the standard 2-cube I^2.

10.2 Day 27: The boundary operator

Goals

By the end of today's class, we should be able to do the following.

1. Determine the **boundary** of a singular k-chain.
2. State some important properties of the boundary operator.

10.2.1 Boundaries

We now define the boundary operator on singular k-chains. We'll see that computing the boundary of a singular k-cube can be somewhat tedious, but not conceptually difficult, since the parameter domain $[0,1]^k$ has an easy-to-describe boundary.

Definition. Fix integers $k \geq 1$, $1 \leq i \leq k$, and $0 \leq a \leq 1$. Then the **$(\mathbf{i},\mathbf{a})$-face** of the standard k-cube $I^k : [0,1]^k \to \mathbb{R}^k$ is the singular $(k-1)$-cube $I^k_{i,a} : [0,1]^{k-1} \to [0,1]^k$ defined by

$$I^k_{(i,a)}(u_1,\ldots,u_{k-1}) := (u_1,\ldots,u_{i-1},\underbrace{a}_{i^{\text{th}}\text{ entry}},u_i,\ldots,u_{k-1}).$$

For any singular k-cube $c : [0,1]^k \to \mathcal{R}$ in $\mathcal{R} \subset \mathbb{R}^n$, the **$(\mathbf{i},\mathbf{a})$-face** of c is the singular $(k-1)$-cube in $\mathcal{R}$ defined by $c_{(i,a)} := c \circ I^k_{(i,a)}$.

Remark. Because a is an integer, we have either $a = 0$ or $a = 1$. So a singular k-cube has $2k$ faces: there are k options for the value of i, and 2 options for the value of a.

The notation $I^k_{(i,a)}$ can be read as follows: we hold the i^{th} component constant, and give it the value a.

Example 10.7. Consider the standard 1-cube $I^1 : [0,1] \to \mathbb{R}$. This cube has two faces: the $(1,0)$-face and the $(1,1)$ face. Each of these is a singular 0-cube, $\{*\} \to [0,1]$, defined by

$$I^1_{(1,0)}(*) = 0 \quad \text{and} \quad I^1_{(1,1)}(*) = 1,$$

respectively. Note that these correspond to the two endpoints of the interval $[0,1]$.

Example 10.8. The standard 2-cube $I^2 : [0,1]^2 \to \mathbb{R}^2$ has four faces, corresponding to the two options for i

and the two options for a. Each of these is a singular 1-cube $[0,1] \to [0,1]^2$, and they are defined by

$$I^2_{(1,0)}(u) = (0,u) \qquad I^2_{(1,1)}(u) = (1,u)$$
$$I^2_{(2,0)}(u) = (u,0) \qquad I^2_{(2,1)}(u) = (u,1).$$

See Figure 10.1. Notice that if $c\colon [0,1]^2 \to \mathcal{R}$ is a not-necessarily-standard singular 2-cube, then the four faces of c are given by

$$c_{(1,0)}(u) = c(0,u) \qquad c_{(1,1)}(u) = c(1,u)$$
$$c_{(2,0)}(u) = c(u,0) \qquad c_{(2,1)}(u) = c(u,1).$$

Again, each is a map $[0,1] \to \mathcal{R}$, where $[0,1]$ has parameter u.

As these examples demonstrate, the boundary of a singular k-cube should be a singular $(k-1)$-chain, a linear combination of the various faces of our k-cube. It may not be clear, however, what coefficient each face should have. A singular cube naturally carries an orientation: for instance, the vertical faces of I^2 (each of which is a singular 1-cube) in Figure 10.1 are oriented bottom-to-top, while the horizontal faces are oriented left-to-right. Notice that the left and top faces of I^2 have orientations which disagree with our usual convention. As a consequence, we give these faces a coefficient of -1 when expressing the boundary ∂I^2.

As a matter of fact, any singular k-cube c induces a boundary orientation on ∂c, and the orientations naturally carried by the faces of c may not always agree with this induced orientation. We now distinguish between positively- and negatively-oriented faces.

Definition. Consider the standard k-cube $I^k\colon [0,1]^k \to \mathbb{R}^k$, for some $k \geq 1$. We say that the face $I^k_{(i,a)}\colon [0,1]^{k-1} \to [0,1]^k$ of I^k is **positively-oriented** (as a face of I^k) if

$$\det\begin{pmatrix} | & | & & | \\ \mathbf{N}_{(i,a)} & \frac{\partial}{\partial u_1}(I^k_{(i,a)}) & \cdots & \frac{\partial}{\partial u_{k-1}}(I^k_{(i,a)}) \\ | & | & & | \end{pmatrix} > 0,$$

where $\mathbf{N}_{(i,a)} := (-1)^{a+1}\mathbf{e}_i \in \mathbb{R}^k$ is the outward-pointing unit normal vector to $I^k_{(i,a)}$. Otherwise, $I^k_{(i,a)}$ is **negatively-oriented** (as a face of I^k). The face $c_{(i,a)}$ of an arbitrary singular k-cube is considered to be **positively-oriented** if $I^k_{(i,a)}$ is so, and otherwise is **negatively-oriented**.

Thankfully, this is not something we have to check very often. Instead, there is a combinatorial test for the orientations of the faces of a singular k-cube.

Exercise 10.9. Show that $c_{(i,a)}$ is positively-oriented if $i+a$ is even, and is negatively-oriented if $i+a$ is odd.

We want the coefficient of $c_{(i,a)}$ to be positive or negative depending on its orientation, so Exercise 10.9 tells us that the coefficient of $c_{(i,a)}$ should be $(-1)^{i+a}$.

Definition. For any integer $k \geq 1$, the **boundary** of a singular k-cube $c\colon [0,1]^k \to \mathcal{R}$ in an open subset $\mathcal{R}$ of $\mathbb{R}^n$ is the singular $(k-1)$-chain ∂c defined by

$$\partial c := \sum_{i=1}^{k}\sum_{a=0}^{1}(-1)^{i+a}c_{(i,a)}.$$

The **boundary** of a singular k-chain $\sigma = \sum_{j=1}^{N} n_j c_j$ is the singular $(k-1)$-chain $\partial\sigma$ defined by

$$\partial\sigma := \sum_{j=1}^{N} n_j\,\partial c_j.$$

The boundary of a singular 0-cube is taken to be 0, as is the boundary of a singular 0-chain[2].

[2]Recall that 0 is the empty linear combination of singular k-cubes. We treat 0 as a singular k-chain for any integer k (perhaps even negative integers).

Example 10.10. Consider the singular 1-cube $c\colon [0,1] \to \mathbb{R}^2$ defined by

$$c(u) := (\cos(2\pi u), \sin(2\pi u)).$$

The image of c is obviously the unit circle in $\mathbb{R}^2$, which we think of as being boundaryless. Let's see what our construction does. First, c has two faces, defined by

$$c_{(1,0)}(*) = c(0) = (1,0) \quad \text{and} \quad c_{(1,1)}(*) = c(1) = (1,0).$$

So in fact the two faces of c are the same parametrized 0-surface. It then follows that

$$\partial c = \sum_{i=1}^{1}\sum_{a=0}^{1}(-1)^{i+a}c_{(i,a)} = -c_{(1,0)} + c_{(1,1)} = 0,$$

and we see that c is indeed boundaryless.

Example 10.11. Consider the singular 2-cube $c\colon [0,1]^2 \to \mathbb{R}^2$ defined by

$$c(u_1,u_2) := ((u_1+1)\cos(2\pi u_2), (u_1+1)\sin(2\pi u_2)),$$

the image of which is the annulus $A = \{(x,y) : 1 \le x^2+y^2 \le 4\}$. As does any 2-cube, c has four faces, defined by

$$\begin{aligned} c_{(1,0)}(u) &= c(0,u) = (\cos(2\pi u), \sin(2\pi u)) \\ c_{(1,1)}(u) &= c(1,u) = (2\cos(2\pi u), 2\sin(2\pi u)) \\ c_{(2,0)}(u) &= c(u,0) = (u+1, 0) \\ c_{(2,1)}(u) &= c(u,1) = (u+1, 0). \end{aligned}$$

Notice that $c_{(2,0)} = c_{(2,1)}$. Draw a picture to convince yourself that this is correct. We then find that

$$\partial c = -c_{(1,0)} + c_{(1,1)} + c_{(2,0)} - c_{(2,1)} = c_{(1,1)} - c_{(1,0)}.$$

Informally, this is telling us that ∂c consists of the outer circle $c_{(1,1)}$ with the counterclockwise orientation, plus the inner circle $c_{(1,0)}$ with the clockwise orientation.

10.2.2 Properties of ∂

For an open subset $\mathcal{R}$ of $\mathbb{R}^n$, we now have two interesting sequences of maps:

$$\Omega^0(\mathcal{R}) \xrightarrow{d} \Omega^1(\mathcal{R}) \xrightarrow{d} \cdots \xrightarrow{d} \Omega^{n-1}(\mathcal{R}) \xrightarrow{d} \Omega^n(\mathcal{R}) \xrightarrow{d} \cdots \tag{10.1}$$

and

$$\mathcal{C}_0(\mathcal{R}) \xleftarrow{\partial} \mathcal{C}_1(\mathcal{R}) \xleftarrow{\partial} \cdots \xleftarrow{\partial} \mathcal{C}_{n-1}(\mathcal{R}) \xleftarrow{\partial} \mathcal{C}_n(\mathcal{R}) \xleftarrow{\partial} \cdots . \tag{10.2}$$

(The sequences continue to make sense when we place larger values in the subscripts or superscripts, but are no longer very interesting.) In this section, we want to show that some of the properties enjoyed by the exterior derivative d are also enjoyed by ∂.

First, recall that if $\Phi\colon \mathcal{D} \to \mathcal{R}$ is a smooth map between open subsets of Euclidean spaces, then we can relate the sequences in (10.1) for the two regions:

$$\begin{array}{ccccccccccc} \Omega^0(\mathcal{R}) & \xrightarrow{d} & \Omega^1(\mathcal{R}) & \xrightarrow{d} & \cdots & \xrightarrow{d} & \Omega^{n-1}(\mathcal{R}) & \xrightarrow{d} & \Omega^n(\mathcal{R}) & \xrightarrow{d} & \cdots \\ \downarrow{\scriptstyle \Phi^*} & & \downarrow{\scriptstyle \Phi^*} & & & & \downarrow{\scriptstyle \Phi^*} & & \downarrow{\scriptstyle \Phi^*} & & \\ \Omega^0(\mathcal{D}) & \xrightarrow{d} & \Omega^1(\mathcal{D}) & \xrightarrow{d} & \cdots & \xrightarrow{d} & \Omega^{n-1}(\mathcal{D}) & \xrightarrow{d} & \Omega^n(\mathcal{D}) & \xrightarrow{d} & \cdots \end{array} . \tag{10.3}$$

Moreover, we know that *the exterior derivative commutes with pullback*, meaning that we can move through this diagram in any order we want: $\Phi^* \circ d = d \circ \Phi^*$. We can also use Φ to relate the sequences in (10.2) for $\mathcal{D}$ and $\mathcal{R}$:

$$\begin{array}{ccccccccccc} \mathcal{C}_0(\mathcal{R}) & \xleftarrow{\partial} & \mathcal{C}_1(\mathcal{R}) & \xleftarrow{\partial} & \cdots & \xleftarrow{\partial} & \mathcal{C}_{n-1}(\mathcal{R}) & \xleftarrow{\partial} & \mathcal{C}_n(\mathcal{R}) & \xleftarrow{\partial} & \cdots \\ \uparrow{\scriptstyle \Phi_*} & & \uparrow{\scriptstyle \Phi_*} & & & & \uparrow{\scriptstyle \Phi_*} & & \uparrow{\scriptstyle \Phi_*} & & \\ \mathcal{C}_0(\mathcal{D}) & \xleftarrow{\partial} & \mathcal{C}_1(\mathcal{D}) & \xleftarrow{\partial} & \cdots & \xleftarrow{\partial} & \mathcal{C}_{n-1}(\mathcal{D}) & \xleftarrow{\partial} & \mathcal{C}_n(\mathcal{D}) & \xleftarrow{\partial} & \cdots \end{array} . \tag{10.4}$$

By analogy with (10.3), we want an equation of the form $\Phi_* \circ \partial = \partial \circ \Phi_*$. That is, we want to say that *the boundary operator commutes with pushforward*. This is indeed the case, and our proof of this fact makes use of the following exercise.

Exercise 10.12. Let $c : [0,1]^k \to \mathcal{R}$ be a singular k-cube in an open subset $\mathcal{R}$ of $\mathbb{R}^n$. Show that $\partial c = c_* \partial I^k$.

Proposition 10.13. *Let $\Phi : \mathcal{D} \to \mathcal{R}$ be a smooth map between open subsets of Euclidean spaces. Then*

$$\partial \Phi_* \sigma = \Phi_* \partial \sigma,$$

for any singular k-chain σ in $\mathcal{D}$.

Proof. First, consider a singular k-cube $c : [0,1]^k \to \mathcal{D}$ in $\mathcal{D}$. Then we have

$$\partial \Phi_* c = \partial(\Phi \circ c) = (\Phi \circ c)_* \partial I^k = \Phi_* c_* \partial I^k = \Phi_* \partial c,$$

where the second and fourth equalities use Exercise 10.12. Next we write $\sigma = \sum_{j=1}^N n_j c_j$ for some singular k-cubes $c_j : [0,1]^k \to \mathcal{D}$. Then we have

$$\partial \Phi_* \sigma = \partial \Phi_* \left(\sum_{j=1}^N n_j c_j \right) = \sum_{j=1}^N n_j \, \partial \Phi_* c_j = \sum_{j=1}^N n_j \, \Phi_* \partial c_j = \Phi_* \partial \left(\sum_{j=1}^N n_j c_j \right) = \Phi_* \partial \sigma,$$

as desired. □

In focusing on the commutativity of (10.3), we overlooked an even more important property of (10.1): the nilpotence of d. We showed in Exercise 9.11 that $d \circ d = 0$, and thus we know that composing consecutive steps in (10.1) always gives the zero map. To use the vocabulary we developed around the exterior derivative, this means that all exact forms (those which are derivatives) are closed (have derivative zero). We define similar (if less creative) vocabulary for the boundary operator.

Definition. We say that a singular k-chain σ is a **boundary** if there is some singular $(k+1)$-chain τ with the property that $\partial \tau = \sigma$. We say that σ is a **cycle** if $\partial \sigma = 0$.

By drawing some pictures, it's probably not hard to convince yourself that every boundary is a cycle. In terms of (10.2), this means that $\partial \circ \partial = 0$, which we now prove.

Theorem 10.14: Nilpotence of ∂

Let $\mathcal{R}$ be an open subset of $\mathbb{R}^n$, and let σ be a singular k-chain in $\mathcal{R}$, for some $k \geq 0$. Then $\partial(\partial \sigma) = 0$.

Proof. We first note that if $k = 0$, then $\partial \sigma = 0$ by definition, and thus $\partial \sigma = 0$. Similarly, if $k = 1$, then $\partial \sigma$ is a 0-chain, and therefore has boundary 0.

We'll prove that if $c : [0,1]^k \to \mathcal{R}$ is a singular k-cube, then $\partial(\partial c) = 0$. The desired result then clearly follows from the linearity of ∂. We have

$$\partial(\partial c) = \partial \left(\sum_{i=1}^k \sum_{a=0}^1 (-1)^{i+a} c_{(i,a)} \right) = \sum_{i=1}^k \sum_{a=0}^1 (-1)^{i+a} \partial c_{(i,a)},$$

with the last equality using the definition of the boundary of a singular $(k-1)$-chain. So we're left to investigate the boundary $\partial c_{(i,a)}$, for some integers $1 \leq i \leq k$ and $0 \leq a \leq 1$. According to the definition of the boundary, we have

$$\partial(c_{(i,a)}) = \sum_{j=1}^{k-1} \sum_{b=0}^1 (-1)^{j+b} (c_{(i,a)})_{(j,b)}.$$

Thus we find that

$$\partial(\partial c) = \sum_{i=1}^k \sum_{a=0}^1 \sum_{j=1}^{k-1} \sum_{b=0}^1 (-1)^{i+a+j+b} (c_{(i,a)})_{(j,b)}$$

With close attention to the indices of our sums, we can do some trickery. For the sake of notation, we will suppress the sums over $0 \leq a, b \leq 1$ for the rest of this argument, but it should be understood that each

expression written is summed over these variables. First, we rewrite:

$$\begin{aligned}\partial(\partial c) &= \sum_{i=1}^{k}\sum_{j=1}^{k-1}(-1)^{i+a+j+b}(c_{(i,a)})_{(j,b)} \\ &= \sum_{i=1}^{k-1}\sum_{j=i}^{k-1}(-1)^{i+a+j+b}(c_{(i,a)})_{(j,b)} + \sum_{i=1}^{k}\sum_{j=1}^{i-1}(-1)^{i+a+j+b}(c_{(i,a)})_{(j,b)}.\end{aligned}$$

Next, we outsource a little bit of work.

Exercise 10.15. Show that if $j < i$, then $(c_{(i,a)})_{(j,b)} = (c_{(j,b)})_{(i-1,a)}$.

We can substitute this into the last sum above:

$$\partial(\partial c) = \sum_{i=1}^{k-1}\sum_{j=i}^{k-1}(-1)^{i+a+j+b}(c_{(i,a)})_{(j,b)} + \sum_{i=1}^{k}\sum_{j=1}^{i-1}(-1)^{i+a+j+b}(c_{(j,b)})_{(i-1,a)}.$$

Now we point out that

$$\sum_{i=1}^{k}\sum_{j=1}^{i-1}(-1)^{i+a+j+b}(c_{(j,b)})_{(i-1,a)} = \sum_{j=1}^{k-1}\sum_{i=j+1}^{k}(-1)^{i+a+j+b}(c_{(j,b)})_{(i-1,a)} = \sum_{j=1}^{k-1}\sum_{i=j}^{k-1}(-1)^{i+1+a+j+b}(c_{(j,b)})_{(i,a)},$$

where the second equality is obtained by replacing the variable i with $i+1$. At last, we swap the roles of i and j, as well as those of a and b, in this last sum:

$$\sum_{j=1}^{k-1}\sum_{i=j}^{k-1}(-1)^{i+1+a+j+b}(c_{(j,b)})_{(i,a)} = \sum_{i=1}^{k-1}\sum_{j=i}^{k-1}(-1)^{j+1+b+i+a}(c_{(i,a)})_{(j,b)}.$$

Substituting this back into our expression for $\partial(\partial c)$ gives

$$\partial(\partial c) = \sum_{i=1}^{k-1}\sum_{j=i}^{k-1}(-1)^{i+a+j+b}(c_{(i,a)})_{(j,b)} + \sum_{i=1}^{k-1}\sum_{j=i}^{k-1}(-1)^{j+1+b+i+a}(c_{(i,a)})_{(j,b)} = 0,$$

as wanted. □

In some ways, the nilpotence of ∂ is more important than the fact that ∂ commutes with pullback; we discussed the properties in this order because we wanted to delay the more tedious proof. The fact that $\partial \circ \partial = 0$ tells us that (10.2) is a **chain complex**, and allows us to feed (10.2) into a gigantic machine called **homological algebra**, the output of which is some information about $\mathcal{R}$. By analogy, we call (10.1) a **cochain complex** (just because d increases the index, rather than decreasing) and feed it into the same machine. The equations

$$\Phi_* \circ \partial = \partial \circ \Phi_* \quad \text{and} \quad \Phi^* \circ d = d \circ \Phi^*$$

then tell us that the pushforward is a **chain map** and the pullback is a **cochain map**, respectively, and properties such as this ensure that continuous changes to $\mathcal{R}$ won't affect the information given to us by the homological algebra machine. That is, these equations somehow ensure that the data we receive from the homological algebra machine is actually data about the *topological type* of $\mathcal{R}$. A truly remarkable fact, known as **De Rham's theorem**, is that the homological algebra machine gives the same output for (10.1) as it does for (10.2) (up to a certain amount of lying). We won't have time to prove this result in our course, but a crucial first step in the proof is Stokes' theorem, which we take up next time.

10.3 Day 28: Stokes' theorem for k-chains

Goals

By the end of today's class, we should be able to do the following.

1. State and prove **Stokes' theorem** for singular chains and differential forms.
2. Interpret the exterior derivative as the natural extension of gradient, curl, or divergence.

The last few weeks may have felt a bit like aimless wandering, as we've defined various abstract notions whose importance may not yet be completely clear. Today we summit our final mountaintop, from which one can see a vast land of manifolds and homology.

To an open subset $\mathcal{R}$ of a Euclidean space $\mathbb{R}^n$, we have now associated two important operations: the boundary and derivative maps

$$\partial : \mathcal{C}_k(\mathcal{R}) \to \mathcal{C}_{k-1}(\mathcal{R}) \quad \text{and} \quad d : \Omega^{k-1}(\mathcal{R}) \to \Omega^k(\mathcal{R}).$$

We've noticed that these two operations enjoy some formally similar properties, such as nilpotence and commutativity with some sort of composition map Φ_* or Φ^*. Our goal today is to prove a crucial relationship between the boundary and derivative operators, from which these similarities follow.

Theorem 10.16: Stokes' theorem for singular k-chains

Fix integers $1 \leq k \leq n$, take an open subset $\mathcal{R}$ of $\mathbb{R}^n$, and let σ be a singular k-chain in $\mathcal{R}$. Then

$$\int_\sigma d\lambda = \int_{\partial\sigma} \lambda,$$

for any differential $(k-1)$ form λ on $\mathcal{R}$.

In terms of the integral pairing $\langle -, - \rangle : \mathcal{C}_k(\mathcal{R}) \times \Omega^k(\mathcal{R}) \to \mathbb{R}$ defined above, Stokes' theorem tells us that ∂ and d are adjoint (as were the pushforward and pullback), because

$$\langle \sigma, d\lambda \rangle = \langle \partial\sigma, \lambda \rangle,$$

for any singular k-chain σ in $\mathcal{R}$ and any $(k-1)$-form λ on $\mathcal{R}$.

This adjunction between ∂ and d is the first in a beautiful sequence of results relating the chain complex $(\mathcal{C}_\bullet(\mathcal{R}), \partial)$ to the cochain complex $(\Omega^\bullet(\mathcal{R}), d)$. These results allow us to show that the geometry we've done in this course can be used to compute *topological* information about the region $\mathcal{R}$, and in fact these results generalize to a setting which is still more abstract. Of course we're out of time to discuss these ideas in this course, but we will see the crucial first step which makes these lines of inquiry possible.

Here on this mountain we have climbed, our course will die; we will see the promised land only from a distance, and will not cross over into it.

10.3.1 The proof of Stokes' theorem

We will prove Theorem 10.16 in stages, considering progressively more general chains and forms. Because a singular k-chain is a linear combination of singular k-cubes, we can reduce to the case of integration over a singular k-cube. But a singular k-cube is simply a pushforward of the standard k-cube, so we consider first the case of integration over the standard k-cube.

The standard k-cube is a singular cube in $\mathbb{R}^k$, so we can write our differential k-form λ as

$$\lambda = \sum_{j=1}^{k} f_j \, dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^k,$$

for some smooth functions $f_1, \ldots, f_k : \mathbb{R}^k \to \mathbb{R}$. We first consider the case where $f_j = 0$ for all but one value of $1 \leq j \leq k$. This will be our only technical result, and the machinery of differential forms will then carry us the rest of the way.

Lemma 10.17. *Fix integers $1 \le j \le k$ and choose a smooth function $f : \mathbb{R}^k \to \mathbb{R}$. Then*

$$\int_{I^k} d(f\, dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^k) = \int_{\partial I^k} f\, dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^k,$$

where I^k is the standard k-cube.

Proof. Let's start by introducing some notation. If we write $I_j = (1, \ldots, \widehat{j}, \ldots, k)$, then we can rewrite our desired equation as

$$\int_{I^k} d(f\, dx^{I_j}) = \int_{\partial I^k} f\, dx^{I_j}. \tag{10.5}$$

The left side of this equation is pretty easy to unwind. One of our first product rules for forms told us that $d(f\, dx^{I_j}) = df \wedge dx^{I_j}$, and since $dx^i \wedge dx^{I_j} = 0$ whenever $i \neq j$, we find that

$$d(f\, dx^{I_j}) = df \wedge dx^{I_j} = \frac{\partial f}{\partial x_j}\, dx^j \wedge dx^{I_j} = (-1)^{j-1} \frac{\partial f}{\partial x_j}\, dx^I,$$

where $I = (1, \ldots, k)$. Moreover, since $\det I^k_* \equiv 1$ and dx^I is a top-degree form on $\mathbb{R}^k$, Theorem 9.21 tells us that

$$(I^k)^*(d(f\, dx^{I_j})) = (I^k)^* \left((-1)^{j-1} \frac{\partial f}{\partial x_j}\, dx^I \right) = (-1)^{j-1} \left(\frac{\partial f}{\partial x_j} \circ I^k \right) du^I.$$

So the left side of our desired equation is given by

$$\int_{I^k} d(f\, dx^{I_j}) = \int_{[0,1]^k} (I^k)^*(d(f\, dx^{I_j})) = (-1)^{j-1} \int_{[0,1]^k} \left(\frac{\partial f}{\partial x_j} \circ I^k \right) du^I = (-1)^{j-1} \int_{[0,1]^k} \frac{\partial}{\partial u_j} (f \circ I^k), \tag{10.6}$$

where the final integral is of the type encountered in part one of the course. In the last step, we used the chain rule to rewrite our function, as well as the definition of integration of top-degree forms on $\mathbb{R}^k$. The right side of 10.5 is somewhat more involved.

Since the boundary ∂I^k is a linear combination of the faces of I^k, we start by considering the integral of $f\, dx^{I_j}$ over each of these faces. That is, we want to investigate

$$\int_{I^k_{(i,a)}} f\, dx^{I_j} = \int_{[0,1]^{k-1}} (I^k_{(i,a)})^*(f\, dx^{I_j}) = \int_{[0,1]^{k-1}} (f \circ I^k_{(i,a)})(I^k_{(i,a)})^*(dx^{I_j}), \tag{10.7}$$

for some integers $1 \le i \le k$ and $0 \le a \le 1$. A good place to start is a computation of $(I^k_{(i,a)})^*(dx^{I_j})$. Notice that this is a top-degree form on $[0,1]^{k-1} \subset \mathbb{R}^{k-1}$, and thus we may write

$$(I^k_{(i,a)})^*(dx^{I_j}) = g\, du^1 \wedge \cdots \wedge du^{k-1},$$

for some smooth function $g : [0,1]^{k-1} \to \mathbb{R}$. We determine g by evaluating this form on the standard basis $\partial_{u_1}, \ldots, \partial_{u_{k-1}}$:

$$(g\, du^1 \wedge \cdots \wedge du^{k-1})(\partial_{u_1}, \ldots, \partial_{u_{k-1}}) = g.$$

Plugging these basis vectors into $(I^k_{(i,a)})^*(dx^{I_j})$, we find that we can write

$$(I^k_{(i,a)})^*(dx^{I_j}) = ((I^k_{(i,a)})^*(dx^{I_j})(\partial_{u_1}, \ldots, \partial_{u_{k-1}}))\, du^1 \wedge \cdots \wedge du^k.$$

Combined with (10.7), this means that

$$\int_{I^k_{(i,a)}} f\, dx^{I_j} = \int_{[0,1]^{k-1}} (f \circ I^k_{(i,a)})\, (I^k_{(i,a)})^*(dx^{I_j})(\partial_{u_1}, \ldots, \partial_{u_{k-1}}), \tag{10.8}$$

where the integral on the right is of the type seen in part one of the course. So we need to compute $(I^k_{(i,a)})^*(dx^{I_j})(\partial_{u_1}, \ldots, \partial_{u_{k-1}})$.

By definition, $(I^k_{(i,a)})^*(dx^{I_j})(\partial_{u_1}, \ldots, \partial_{u_{k-1}}) = dx^{I_j}((I^k_{(i,a)})_* \partial_{u_1}, \ldots, (I^k_{(i,a)})_* \partial_{u_{k-1}})$, so let's think about the Jacobian matrix $(I^k_{(i,a)})_*$. We have

$$I^k_{(i,a)}(u_1, \ldots, u_{k-1}) = (u_1, \ldots, u_{i-1}, \underbrace{a}_{i^{\text{th}} \text{ entry}}, u_i, \ldots, u_{k-1}).$$

So if $1 \le \ell \le i-1$, then only the ℓ^{th} component of $I^k_{(i,a)}$ has a nonzero derivative with respect to u_ℓ, and this derivative is 1. That is,

$$(I^k_{(i,a)})_* \partial_{u_\ell} = \partial_{x_\ell}, \quad \text{if } 1 \le \ell \le i-1.$$

For $i \le \ell \le k-1$, it is the $(\ell+1)^{\text{st}}$ component of $I^k_{(i,a)}$ which varies with respect to u_ℓ, and we find that

$$(I^k_{(i,a)})_* \partial_{u_\ell} = \partial_{x_{\ell+1}}, \quad \text{if } i \le \ell \le k-1.$$

Putting things together,

$$\begin{aligned}(I^k_{(i,a)})^*(dx^{I_j})(\partial_{u_1},\dots,\partial_{u_{k-1}}) &= dx^{I_j}((I^k_{(i,a)})_*\partial_{u_1},\dots,(I^k_{(i,a)})_*\partial_{u_{k-1}}) = dx^{I_j}(\partial_{x_1},\dots,\widehat{\partial_{x_i}},\dots,\partial_{x_k}) \\ &= \begin{cases} 1, & \text{if } i = j \\ 0, & \text{if } i \ne j \end{cases} =: \delta_{ij}.\end{aligned}$$

This last notation is called the *Kronecker delta*, and is defined by the equation given here. It's just a useful bit of notation.

Substituting this computation into (10.8), leads us to

$$\int_{I^k_{(i,a)}} f\, dx^{I_j} = \delta_{ij} \int_{[0,1]^{k-1}} f \circ I^k_{(i,a)}. \tag{10.9}$$

At last we can compute the right side of (10.5). We have

$$\begin{aligned}\int_{\partial I^k} f\, dx^{I_j} &= \sum_{i=1}^{k}\sum_{a=0}^{1}(-1)^{i+a}\left(\int_{I^k_{(i,a)}} f\, dx^{I_j}\right) = \sum_{i=1}^{k}\sum_{a=0}^{1}(-1)^{i+a}\left(\delta_{ij}\int_{[0,1]^{k-1}} f\circ I^k_{(i,a)}\right) \\ &= (-1)^j \sum_{a=0}^{1}(-1)^a\left(\int_{[0,1]^{k-1}} f\circ I^k_{(j,a)}\right) = (-1)^{j+1}\int_{[0,1]^{k-1}}(f\circ I^k_{(j,1)} - f\circ I^k_{(j,0)}).\end{aligned} \tag{10.10}$$

We emphasize once again that this last integral is of the type seen in part one of the course; namely, it's an integral of a function.

Finally, it's time to compare the left and right hand sides of our desired equation. We expressed the left side as an integral over $[0,1]^k$ in (10.6). Using Fubini's theorem, we have

$$\int_{I^k} d(f\, dx^{I_j}) = (-1)^{j-1}\int_{[0,1]^k}\frac{\partial}{\partial u_j}(f\circ I^k) = (-1)^{j-1}\int_{[0,1]^{k-1}}\left(\int_{[0,1]}\frac{\partial}{\partial u_j}(f\circ I^k)\right), \tag{10.11}$$

where the integral over $[0,1]$ uses the variable u_j, and the outer integral uses $u_1,\dots,\widehat{u_j},\dots,u_k$. With all the indices we've had to track, this entire proof probably feels pretty technical; but here comes the key technical step: we apply the single-variable Fundamental Theorem of Integral Calculus to the integral over $[0,1]$ to see that

$$\begin{aligned}\int_{[0,1]}\frac{\partial}{\partial u_j}(f\circ I^k) &= (f\circ I^k)(u_1,\dots,\overbrace{1}^{j^{\text{th}}\text{ entry}},\dots,u_k) - (f\circ I^k)(u_1,\dots,\overbrace{0}^{j^{\text{th}}\text{ entry}},\dots,u_k) \\ &= (f\circ I^k_{(j,1)})(u_1,\dots,\widehat{u_j},\dots,u_k) - (f\circ I^k_{(j,0)})(u_1,\dots,\widehat{u_j},\dots,u_k).\end{aligned}$$

We're almost there. Substituting this last observation into (10.11), we find that

$$\int_{I^k} d(f\, dx^{I_j}) = (-1)^{j-1}\int_{[0,1]^{k-1}}\left(f\circ I^k_{(j,1)} - f\circ I^k_{(j,0)}\right) = \int_{\partial I^k} f\, dx^{I_j},$$

where this last equality uses (10.10), as well as the fact that $(-1)^{j-1} = (-1)^{j+1}$. But this is precisely our desired equation! So Stokes' theorem holds in this special case. □

With this basic version of Stokes' theorem in hand, we obtain the result for more general forms by using some basic properties of forms.

Lemma 10.18. *For any $k \geq 1$ and any differential $(k-1)$-form λ on $\mathbb{R}^k$,*

$$\int_{I^k} d\lambda = \int_{\partial I^k} \lambda,$$

where I^k is the standard k-cube.

Proof. Because λ is a $(k-1)$-form on $\mathbb{R}^k$, we can write

$$\lambda = \sum_{j=1}^{k} f_j \, dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^k,$$

for some smooth functions $f_1, \ldots, f_k : \mathbb{R}^k \to \mathbb{R}$. Then we find that

$$\begin{aligned}
\int_{I^k} d\lambda &= \int_{I^k} d\left(\sum_{j=1}^{k} f_j \, dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^k \right) = \int_{I^k} \left(\sum_{j=1}^{k} d(f_j \, dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^k) \right) \\
&= \sum_{j=1}^{k} \left(\int_{I^k} d(f_j \, dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^k) \right) = \sum_{j=1}^{k} \left(\int_{\partial I^k} f_j \, dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^k \right) \\
&= \int_{\partial I^k} \left(\sum_{j=1}^{k} f_j \, dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^k \right) = \int_{\partial I^k} \lambda,
\end{aligned}$$

with the antepenultimate equality following from Lemma 10.17. □

So we now know that Stokes' theorem holds for any $(k-1)$-form λ, provided the chain over which we're integrating is the standard k-cube. Let's enhance this to an arbitrary singular k-cube.

Lemma 10.19. *Fix integers $1 \leq k \leq n$, take an open subset $\mathcal{R}$ of $\mathbb{R}^n$, and let $c : [0,1]^k \to \mathcal{R}$ be a singular k-cube in $\mathcal{R}$. Then*

$$\int_c d\lambda = \int_{\partial c} \lambda,$$

for any differential $(k-1)$-form λ on $\mathcal{R}$.

Proof. Recall that we can realize the k-cube c as the pushforward under the map $c : [0,1]^k \to \mathcal{R}$ of the standard k-cube $I^k : [0,1]^k \to [0,1]^k$. Then, using Theorem 10.6, we find that

$$\begin{aligned}
\int_c d\lambda &= \int_{c_* I^k} d\lambda = \int_{I^k} c^*(d\lambda) = \int_{I^k} d(c^*\lambda) \\
&= \int_{\partial I^k} c^*\lambda = \int_{c_* \partial I^k} \lambda = \int_{\partial c} \lambda.
\end{aligned}$$

This computation uses Lemma 10.18 in its fourth step, in addition to using the fact that pushforward commutes with the boundary operator. □

Finally, obtaining Theorem 10.16 from Lemma 10.19 simply uses the linearity of integration over chains.

Proof of Theorem 10.16. Since σ is a singular k-chain, we can write

$$\sigma = \sum_{j=1}^{N} n_j \, c_j$$

for some finite collection of integers $n_j \in \mathbb{Z}$ and singular k-cubes c_j. Then, using Lemma 10.19 and the definition of integration over chains, we have

$$\int_\sigma d\lambda = \sum_{j=1}^{N} n_j \int_{c_j} d\lambda = \sum_{j=1}^{N} n_j \int_{\partial c_j} \lambda = \int_{\partial \sigma} \lambda,$$

as desired. □

Notice that Lemma 10.17 was the only step with any technical content, where we used the single-variable FTIC. The remaining steps were routine applications of standard properties for differential forms and singular chains.

10.3.2 An example

Let's see Stokes' theorem in action.

Example 10.20. Consider the upper hemisphere $\mathcal{S} \subset \mathbb{R}^4$ of the unit 3-sphere S^3 in $\mathbb{R}^4$. That is,

$$\mathcal{S} := \{(x_1, x_2, x_3, x_4) : x_1^2 + x_2^2 + x_3^2 + x_4^2 = 1 \text{ and } x_4 \geq 0\}.$$

Suppose we have a singular 3-cube $c : [0,1]^3 \to \mathbb{R}^4$ which is injective on $(0,1)^3$, and whose image is $\mathcal{S}$. Suppose further that

$$\det \begin{pmatrix} | & | & | & | \\ \mathbf{N} & \frac{\partial c}{\partial u_1} & \frac{\partial c}{\partial u_2} & \frac{\partial c}{\partial u_3} \\ | & | & | & | \end{pmatrix} > 0,$$

where $\mathbf{N}$ is the outward-pointing unit normal vector to $\mathcal{S}$ in $\mathbb{R}^4$. Let's compute $\int_{\mathcal{S}} \eta := \int_c \eta$, where

$$\eta := \sum_{j=1}^{4} (1 - x_j)\, dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^4.$$

At this point, it's hard to be too hopeful about applying Stokes' theorem, since η is not presented as a derivative and $\mathcal{S}$ is not presented as a boundary. Notice, however, that

$$\begin{aligned} d\eta &= \sum_{j=1}^{4} \frac{\partial}{\partial x_j}(1 - x_j)\, dx^j \wedge dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^4 \\ &= \sum_{j=1}^{4} -1\, dx^j \wedge dx^1 \wedge \cdots \wedge \widehat{dx^j} \wedge \cdots \wedge dx^4 \\ &= \sum_{j=1}^{4} (-1)^j\, dx^I = 0. \end{aligned}$$

So if we could realize $\mathcal{S}$ as (part of) a boundary in $\mathbb{R}^4$, then we could trade this boundary in for a derivative, which would kill the form we're integrating. For instance, consider the region

$$\mathcal{W} := \{(x_1, x_2, x_3, x_4) : x_1^2 + x_2^2 + x_3^2 + x_4^2 \leq 1 \text{ and } x_4 \geq 0\}.$$

As a region in $\mathbb{R}^4$, we can think of the boundary $\mathcal{W}$ as consisting of two pieces:

$$\begin{aligned} \partial \mathcal{W} = \{&(x_1, x_2, x_3, x_4) : x_1^2 + x_2^2 + x_3^2 + x_4^2 = 1 \text{ and } x_4 \geq 0\} \\ &\cup \{(x_1, x_2, x_3, x_4) : x_1^2 + x_2^2 + x_3^2 + x_4^2 \leq 1 \text{ and } x_4 = 0\}. \end{aligned}$$

The first of these pieces is our hemisphere $\mathcal{S}$, while the second is the 3-ball in $\mathbb{R}^3$ of radius 1, included into $\mathbb{R}^4$ in the obvious way. We'll call this second piece $\mathcal{B}$. Of course we're being a bit sloppy here, since we haven't expressed any of our regions or hypersurfaces as parametrized k-surfaces, but the point is that this general Stokes' theorem continues to work as expected. Namely, we have

$$\int_{\mathcal{W}} d\eta = \int_{\partial \mathcal{W}} \eta = \int_{\mathcal{S}} \eta + \int_{\mathcal{B}} \eta,$$

with the hypersurfaces on the right carrying the boundary orientation. Since $d\eta = 0$, the integral on the left is 0, so we have

$$\int_{\mathcal{S}} \eta = -\int_{\mathcal{B}} \eta,$$

with $\mathcal{S}$ oriented by an outward-pointing normal vector (as we required above), and $\mathcal{B}$ oriented by $-\partial_{x_4}$. Now on $\mathcal{B}$ we have $x_4 \equiv 0$ and $dx^4 \equiv 0$, meaning that

$$\eta|_{\mathcal{B}} = dx^1 \wedge dx^2 \wedge dx^3.$$

Consider the singular 3-cube $c_{\mathcal{B}} : [0,1]^3 \to \mathbb{R}^4$ defined by

$$c_{\mathcal{B}}(u_1, u_2, u_3) := (u_1 \sin(\pi u_2) \cos(2\pi u_3), u_1 \sin(\pi u_2) \sin(2\pi u_3), u_1 \cos(\pi u_2), 0).$$

Inspired by spherical coordinates, this 3-cube is injective on the interior of its parameter domain, and its image is $\mathcal{B}$. Moreover, since

$$\begin{aligned}
&\det\begin{pmatrix} | & | & | & | \\ -\partial_{x_4} & \frac{\partial c_{\mathcal{B}}}{\partial u_1} & \frac{\partial c_{\mathcal{B}}}{\partial u_2} & \frac{\partial c_{\mathcal{B}}}{\partial u_3} \\ | & | & | & | \end{pmatrix} \\
&\qquad = \det\begin{pmatrix} 0 & \sin(\pi u_2)\cos(2\pi u_3) & \pi u_1 \cos(\pi u_2)\cos(2\pi u_3) & -2\pi u_1 \sin(\pi u_2)\sin(2\pi u_3) \\ 0 & \sin(\pi u_2)\sin(2\pi u_3) & \pi u_1 \cos(\pi u_2)\sin(2\pi u_3) & 2\pi u_1 \sin(\pi u_2)\cos(2\pi u_3) \\ 0 & \cos(\pi u_2) & -\pi u_1 \sin(\pi u_2) & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix} \\
&\qquad = 2\pi^2 u_1^2 \sin(\pi u_2) > 0,
\end{aligned}$$

for all $(u_1, u_2, u_3) > 0$, we see that $c_{\mathcal{B}}$ is positively-oriented as a face of $\partial\mathcal{W}$. So we can integrate η over $\mathcal{B}$ by integrating $c_{\mathcal{B}}^*\eta$ over $[0,1]^3$. Let's consider $c_{\mathcal{B}}^*\eta$.

First, notice that $c_{\mathcal{B}}^*\eta$ is a top-degree form on $[0,1]^3$. This means that we can write

$$c_{\mathcal{B}}^*\eta = (c_{\mathcal{B}}^*\eta)(\partial_{u_1}, \partial_{u_2}, \partial_{u_3})\, du^1 \wedge du^2 \wedge du^3,$$

so we need to compute $(c_{\mathcal{B}}^*\eta)(\partial_{u_1}, \partial_{u_2}, \partial_{u_3})$. Since $\eta|_{\mathcal{B}} = dx^1 \wedge dx^2 \wedge dx^3$, we see that

$$\begin{aligned}
(c_{\mathcal{B}}^*\eta)(\partial_{u_1}, \partial_{u_2}, \partial_{u_3}) &= (dx^1 \wedge dx^2 \wedge dx^3)((c_{\mathcal{B}})_*\partial_{u_1}, (c_{\mathcal{B}})_*\partial_{u_2}, (c_{\mathcal{B}})_*\partial_{u_3}) \\
&= \det\begin{pmatrix} | & | & | \\ (c_{\mathcal{B}})_*\partial_{u_1} & (c_{\mathcal{B}})_*\partial_{u_2} & (c_{\mathcal{B}})_*\partial_{u_3} \\ | & | & | \end{pmatrix}_{(1,2,3)} \\
&= \det\begin{pmatrix} \sin(\pi u_2)\cos(2\pi u_3) & \pi u_1 \cos(\pi u_2)\cos(2\pi u_3) & -2\pi u_1 \sin(\pi u_2)\sin(2\pi u_3) \\ \sin(\pi u_2)\sin(2\pi u_3) & \pi u_1 \cos(\pi u_2)\sin(2\pi u_3) & 2\pi u_1 \sin(\pi u_2)\cos(2\pi u_3) \\ \cos(\pi u_2) & -\pi u_1 \sin(\pi u_2) & 0 \end{pmatrix} \\
&= 2\pi^2 u_1^2 \sin(\pi u_2).
\end{aligned}$$

So at last we have

$$\int_{\mathcal{B}} \eta = \int_{c_{\mathcal{B}}} \eta = \int_{[0,1]^3} c_{\mathcal{B}}^*\eta = \int_{[0,1]^3} 2\pi^2 u_1^2 \sin(\pi u_2)\, du^1 \wedge du^2 \wedge du^3 = \int_{[0,1]^3} 2\pi^2 u_1^2 \sin(\pi u_2),$$

where the last integral is of the type seen in part one of the course. Hopefully we recognize this last integral as computing the volume of the unit 3-ball in $\mathbb{R}^3$. Whether we do or not, we can verify that this last integral gives us $4\pi/3$, and thus

$$\int_{\mathcal{S}} \eta = -\int_{\mathcal{B}} \eta = \boxed{-\frac{4\pi}{3}.}$$

Whew! Hopefully we've made Stokes proud.

10.3.3 Reinterpreting the exterior derivative

Recall that in Sections 5.3.3 and 6.3.3 we used an FTIC to give the derivatives curl and div some geometric meaning. We can now use Stokes' theorem to reinterpret the exterior derivative of a differential k-form. If one has the perspective that a differential form is "a thing to be integrated," then this is perhaps the *correct* definition of the derivative.

Throughout this section, we'll suppose that we have a differential k-form λ on $\mathbb{R}^n$, but that we don't know how to compute the $(k+1)$-form $d\lambda$. We will *assume* Stokes' theorem, and compute $d\lambda$ using this assumption.

Crucial to our earlier discussions was the fact that if $f : \mathbb{R} \to \mathbb{R}$ is a continuous function, then

$$\lim_{r\to 0}\left(\frac{1}{2r}\int_{x_0-r}^{x_0+r} f(x)\, dx\right) = f(x_0) \tag{10.12}$$

for any point $x_0 \in \mathbb{R}$. You can prove this fact using the usual Fundamental Theorem of Integral Calculus, but we claim that this is an equation which makes good intuitive sense. The expression in parentheses on the

left computes the average value of f over the interval $[x_0 - r, x_0 + r]$; as r shrinks to 0, so that this interval shrinks down to the point $\{x_0\}$, the average value of f over the interval should tend towards $f(x_0)$, provided f is continuous.

A strange way of interpreting the above discussion is as follows: a continuous function $f : \mathbb{R} \to \mathbb{R}$ is completely determined by its average values of the various closed intervals of $\mathbb{R}$. That is, if we know the average value of f over every closed interval I of nonzero length, then we can use (10.12) to determine every value of f. We'll use this interpretation to compute $d\lambda$: using Stokes' theorem and integration of λ, we know the "average value" of $d\lambda$ over any singular $(k+1)$-chain in $\mathbb{R}^n$, and we can use these values to determine $d\lambda$.

The differential form $d\lambda$ that we're trying to compute is a $(k+1)$-form, and thus may be written as

$$d\lambda = \sum_I f_I \, dx^I,$$

for some smooth functions $f_I : \mathbb{R}^n \to \mathbb{R}$, where the sum is taken over standard multi-indices of length $k+1$. Our job, then, is to determine the functions f_I. Notice that if $I = (i_1, \ldots, i_{k+1})$, then

$$f_I = d\lambda(\partial_{i_1}, \ldots, \partial_{i_{k+1}}),$$

since $dx^J(\partial_{i_1}, \ldots, \partial_{i_{k+1}}) = 0$ if $J \neq I$. Here we're writing ∂_j for the coordinate vector field ∂_{x_j}. So we can determine $f_I(p)$ by computing $(d\lambda)_p(\partial_{i_1}, \ldots, \partial_{i_{k+1}})$, for any point $p \in \mathbb{R}^n$. We claim that we can approximate the value $(d\lambda)_p(\partial_{i_1}, \ldots, \partial_{i_{k+1}})$ with the average value of $d\lambda$ over a singular $(k+1)$-cube whose (small) image contains p. For each $r \in \mathbb{R}$ we define a singular $(k+1)$-cube $c_r : [0,1]^{k+1} \to \mathbb{R}^n$ by

$$c_r(u_1, \ldots, u_{k+1}) := p + r\,u_1\,\partial_{i_1} + \cdots + r\,u_{k+1}\,\partial_{i_{k+1}}.$$

The image of c_r is the $(k+1)$-dimensional parallelotope based at p and spanned by $r\,\partial_{i_1}, \ldots, r\,\partial_{i_{k+1}}$, and if $|r| > 0$ is very small, we expect $f_I(p)$ to be approximated by the average value of $d\lambda$ over c_r. That is, we expect

$$f_I(p) \approx \frac{1}{\mathrm{vol}(c_r)} \int_{c_r} d\lambda = \frac{1}{r^{k+1}} \int_{c_r} d\lambda,$$

where $\mathrm{vol}(c_r)$ is the $(k+1)$-dimensional volume of c_r, equal to r^{k+1} because the vectors $\partial_{i_1}, \ldots, \partial_{i_{k+1}}$ are orthonormal in $\mathbb{R}^n$.

We can now invoke Stokes' theorem to rewrite the integral of $d\lambda$ (the form we are trying to define) over c_r as the integral of λ over ∂c_r. We thus find that

$$f_I(p) = \lim_{r \to 0} \left(\frac{1}{r^{k+1}} \int_{\partial c_r} \lambda \right).$$

From here, we may extend by linearity to find that if p is any point in $\mathbb{R}^n$ and $\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}$ is any vector based at p, then we can define a singular $(k+1)$-cube

$$V_r(u_1, \ldots, u_{k+1}) := p + r\,u_1\,\mathbf{v}_1 + \cdots + r\,u_{k+1}\,\mathbf{v}_{k+1}$$

in $\mathbb{R}^n$ for any $r \in \mathbb{R}$, and we will have

$$(d\lambda)_p(\mathbf{v}_1, \ldots, \mathbf{v}_{k+1}) = \lim_{r \to 0} \left(\frac{1}{r^{k+1}} \int_{\partial V_r} \lambda \right). \tag{10.13}$$

For those who define differential forms as *things which are to be integrated*, (10.13) is probably the correct way to define the derivative of a differential form. However we define the derivative, we want Stokes' theorem to hold, and (10.13) ensures that this result holds on an infinitesimal level.

Just in case you find (10.13) to be less-than-natural, let's consider the case $k = 0$ — that is, the case of a smooth function f on $\mathbb{R}^n$. In this case, df accepts as input a point p and a vector $\mathbf{v}$. Equation 10.13 tells us to consider the integral of f over the boundary of the singular 1-cube

$$V_r(u) = p + r\,u\,\mathbf{v}.$$

Notice that ∂V_r consists of two faces:

$$(\partial V_r)_{(1,1)} : * \mapsto p + r\,\mathbf{v} \quad \text{and} \quad (\partial V_r)_{(1,0)} : * \mapsto p,$$

the first with positive coefficient and the second with negative. We thus find that

$$\int_{\partial V_r} f = \int_{(\partial V_r)_{(1,1)}} f - \int_{(\partial V_r)_{(1,0)}} f = f(p + r\,\mathbf{v}) - f(p).$$

So (10.13) tells us that

$$df_p(\mathbf{v}) = \lim_{r\to 0}\left(\frac{1}{r}\int_{\partial V_r} f\right) = \lim_{r\to 0}\frac{f(p + r\,\mathbf{v}) - f(p)}{r}.$$

This matches the first definition we saw for directional derivatives, and suggests that (10.13) generalizes the familiar notion of derivative.