

Math 1610 · Honors Calculus I · Fall 2026

Grad and Curl with Matrices and a Bit of Linear Algebra

Danny Krashen · torsorlab

Contents

The question 4

1 A bit of linear algebra 7

2 The derivative of a function: a row, or a vector 19

3 The derivative of a vector field: a matrix 21

4 Curl: the antisymmetric part 24

5 Gradients have zero curl 28

6 Closed fields on a disk are gradients 31

7 In three dimensions: the curl vector 36

When you come to a fork in the road, take it.

— Yogi Berra

There is no there there.

— Gertrude Stein

I told the doctor I broke my leg in two places. He told me to quit going to those places.

— Henny Youngman

The question

Arrows on the plane:
is there a hill they climb from?
Ask the tiny wheel.

Here is a vector field $\vec{F}$. Is it the gradient of something?

Here is the answer these notes arrive at. Compute the *curl* of $\vec{F}$, which is the antisymmetric part of its matrix of partial derivatives (doubled). If $\vec{F}$ is a gradient, then its curl is zero. Conversely, on a disk, if the curl of $\vec{F}$ is zero, then $\vec{F}$ is a gradient.

We care because gradients are the vector fields whose line integrals obey a version of the second fundamental theorem of calculus. Recall that theorem in one variable: if G is differentiable on $[a, b]$ and its derivative G' is continuous, then

$$\int_a^b G'(t)\,dt \;=\; G(b) - G(a).$$

To integrate a derivative, we only need to know the original function at the two endpoints. Line integrals of gradients work the same way. If $\vec{F} = \nabla G$, where G has continuous partial derivatives, then for every curve C,

$$\int_C \nabla G \cdot d\vec{r} \;=\; G(\text{end of } C) - G(\text{start of } C).$$

Item by item, the comparison is this: the derivative G' becomes the gradient ∇G; the interval $[a, b]$ becomes the curve C; its two endpoints a and b become the two endpoints of C; and in both cases the integral of a derivative is the difference of the function's values at the ends. This is a gift of almost suspicious generosity, because the shape of the curve stops mattering. The curve can wander through

three counties and a parking garage, and the integral only remembers where it began and where it ended. Work done by gravity is like this. Work done against friction is not, which is why pushing a couch around the living room and back to where it started still leaves you tired.

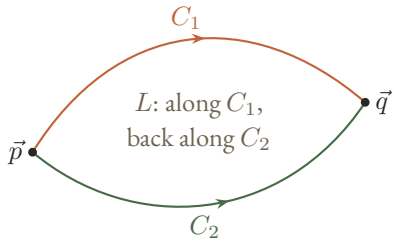
So, given a vector field $\vec{F}$, we would like to know whether we are in this lucky case. A first idea is to look at *closed loops*: curves that end at the same point where they start.

If $\vec{F} = \nabla G$ and C is a closed loop, the formula above gives $G(\text{end}) - G(\text{start}) = 0$, because the end and the start are the same point. So the integral of a gradient around every closed loop is zero.

Closed loops also give a way to test whether line integrals depend only on the endpoints. Suppose C_1 and C_2 are two curves from a point $\vec{p}$ to a point $\vec{q}$. Follow C_1 from $\vec{p}$ to $\vec{q}$, and then follow C_2 backwards, from $\vec{q}$ back to $\vec{p}$. Together these make a closed loop L, in two pieces. Running a curve backwards changes the sign of its line integral, so

$$\int_L \vec{F} \cdot d\vec{r} = \int_{C_1} \vec{F} \cdot d\vec{r} - \int_{C_2} \vec{F} \cdot d\vec{r}.$$

So the integral around L is zero exactly when the integrals along C_1 and C_2 are equal. In other words, "the integral around every closed loop is zero" and "line integrals depend only on the endpoints" are the same condition, seen from two sides.



It is also true, though not proved in these notes, that this condition is enough to make $\vec{F}$ a gradient, at least for a continuous field on an open set in which any two points can be joined by a curve. The idea is to fix a starting point $\vec{p}$ and to define $G(\vec{x})$ as the integral of $\vec{F}$ along any curve from $\vec{p}$ to $\vec{x}$; the condition says

that the choice of curve does not matter.

So one method for deciding whether $\vec{F}$ is a gradient is to integrate $\vec{F}$ around every closed loop and check that all the answers are zero. There are rather a lot of closed loops in the plane, so a faster method would be welcome.

These notes find one. Here is the plan.

1. The derivative of a function $\mathbb{R}^n \to \mathbb{R}$ at a point is a row of numbers (a covector), or equivalently a vector (the gradient).
2. The derivative of a vector field $\vec{F}\colon \mathbb{R}^n \to \mathbb{R}^n$ at a point is a square matrix, written $D\vec{F}$.
3. Every square matrix splits into a symmetric part and an antisymmetric part. The antisymmetric part of $D\vec{F}$, doubled, is called the *curl* of $\vec{F}$.
4. If $\vec{F}$ is the gradient of a function, then the curl of $\vec{F}$ is zero. This turns out to be the statement that mixed partial derivatives can be taken in either order, in disguise.
5. On a disk, the converse is true: if the curl of $\vec{F}$ is zero, then $\vec{F}$ is the gradient of some function.

The words "on a disk" in step 5 are doing real work. If the region has a hole in it, a field can have curl zero everywhere and still fail to be a gradient. Chapter 6 gives an example and explains why the hole is to blame.

Step 2 needs matrices, which you have seen before but which may have faded a little. So we begin there.

Throughout, $\vec{e}_1, \ldots, \vec{e}_n$ are the standard basis vectors of $\mathbb{R}^n$: $\vec{e}_i$ has a 1 in position i and 0 everywhere else.

Chapter 1

A bit of linear algebra

A linear map, so they say,
Is fixed once the $\vec{e}$'s find their way.
Write them down in a row—
No, in columns, just so—
And it splits into stretch plus a sway.

The derivative of a vector field will be a matrix, and its curl will be the antisymmetric part of that matrix. This chapter collects the linear algebra needed for that.

1.1 Linear maps are determined by very little

Definition 1.1. A function $T\colon \mathbb{R}^n \to \mathbb{R}^m$ is *linear* if $T(\vec{v}+\vec{w}) = T(\vec{v})+T(\vec{w})$ and $T(c\vec{v}) = cT(\vec{v})$ for all vectors $\vec{v}, \vec{w} \in \mathbb{R}^n$ and all numbers c.

Linear maps are the well-behaved citizens of the function world. They send lines to lines, they send the origin to the origin, and they never do anything surprising when you are not looking. Their best feature is that they are cheap to describe. Every $\vec{v} \in \mathbb{R}^n$ is $\vec{v} = v_1\vec{e}_1 + \cdots + v_n\vec{e}_n$, so linearity gives

$$T(\vec{v}) = v_1 T(\vec{e}_1) + \cdots + v_n T(\vec{e}_n).$$

Once you know where the n basis vectors go, you know where *every* vector goes. A function with infinitely many inputs is pinned down by just n of its values. Bureaucracies should be so efficient.

Column notation. So far we have written vectors as ordered tuples, like $\vec{v} = (v_1, \ldots, v_n)$. When matrices are involved, it is more convenient to write the

same vector as a *column*:

$$\vec{v} = \begin{pmatrix} v_1 \\ \vdots \\ v_n \end{pmatrix}, \qquad \text{for example} \qquad (2,1) = \begin{pmatrix} 2 \\ 1 \end{pmatrix}.$$

This is only a change of notation, not a new kind of object: the tuple and the column list the same entries in the same order. In running text we will still write tuples, because they fit on a line, but whenever a vector meets a matrix it is a column. (Rows of numbers will get a different job below.)

We record the n vectors $T(\vec{e}_1), \dots, T(\vec{e}_n)$, written as columns, side by side in an $m \times n$ array, the *matrix* of T:

$$A = \begin{pmatrix} | & & | \\ T(\vec{e}_1) & \cdots & T(\vec{e}_n) \\ | & & | \end{pmatrix}, \qquad a_{ij} = \text{the } i\text{-th entry of } T(\vec{e}_j).$$

We write $T(\vec{v}) = A\vec{v}$. The i-th entry of $A\vec{v}$ is

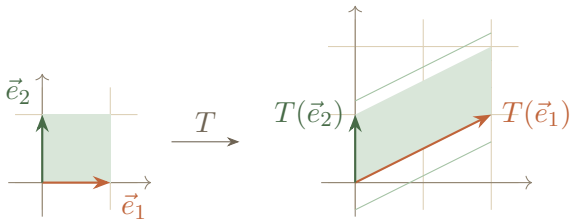
$$(A\vec{v})_i = a_{i1}v_1 + \cdots + a_{in}v_n = (\text{row } i \text{ of } A) \cdot \vec{v}.$$

So there are two ways to read a matrix, and both are useful. Read by columns, it says where the basis vectors go. Read by rows, it is a stack of dot products waiting to happen.

Example 1.2. Let $T\colon \mathbb{R}^2 \to \mathbb{R}^2$ send $\vec{e}_1 \mapsto (2,1)$ and $\vec{e}_2 \mapsto (0,1)$. Then

$$A = \begin{pmatrix} 2 & 0 \\ 1 & 1 \end{pmatrix}, \qquad T(3,-1) = 3\begin{pmatrix} 2 \\ 1 \end{pmatrix} - \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 6 \\ 2 \end{pmatrix}.$$

The unit square, with corners $\vec{0}, \vec{e}_1, \vec{e}_1 + \vec{e}_2, \vec{e}_2$, goes to the parallelogram with corners $\vec{0}, (2,1), (2,2), (0,1)$. Grid lines go to grid lines, slanted. That is what "linear" looks like.



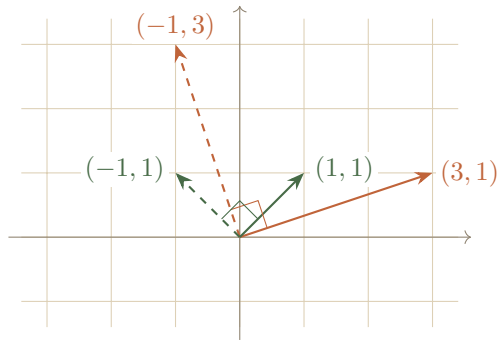
Example 1.3. Consider the linear map $(x, y) \mapsto (-y, x)$. It sends $\vec{e}_1 = (1, 0) \mapsto (0, 1) = \vec{e}_2$ and $\vec{e}_2 = (0, 1) \mapsto (-1, 0) = -\vec{e}_1$, so its matrix is

$$J = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}.$$

Here is what it does to a few vectors:

$$J\begin{pmatrix} 3 \\ 1 \end{pmatrix} = \begin{pmatrix} -1 \\ 3 \end{pmatrix}, \qquad J\begin{pmatrix} 1 \\ 1 \end{pmatrix} = \begin{pmatrix} -1 \\ 1 \end{pmatrix}, \qquad J\begin{pmatrix} -1 \\ 3 \end{pmatrix} = \begin{pmatrix} -3 \\ -1 \end{pmatrix}.$$

Plotted, each output looks like the input turned a quarter turn counterclockwise:



The pictures can be checked by computation, in three ways.

- *Lengths are unchanged.* The vector $(-y, x)$ has length $\sqrt{y^2 + x^2}$, the same as (x, y). For instance, $(3, 1)$ and $(-1, 3)$ both have length $\sqrt{10}$.
- *The angle is a right angle.* The dot product of (x, y) with $(-y, x)$ is $-xy + yx = 0$, so the output is always perpendicular to the input.
- *The turn is counterclockwise.* Write (x, y) in polar form, $(r\cos\theta,\ r\sin\theta)$.

Then

$$(-y, x) = (-r\sin\theta,\ r\cos\theta) = \big(r\cos(\theta+\tfrac{\pi}{2}),\ r\sin(\theta+\tfrac{\pi}{2})\big),$$

by the addition formulas for sine and cosine. The length stays r and the angle goes up by $\frac{\pi}{2}$: a rotation by $90°$ counterclockwise. The first two bullets are this fact seen without polar coordinates.

So J is the *quarter turn*. Doing it twice should be a half turn, which sends every vector to its negative, and indeed $J(J(x,y)) = J(-y,x) = (-x,-y)$; the third computation above is an instance, since $(-1,3)$ was itself $J(3,1)$. It is worth remembering J. It will keep turning up, which is in its nature.

Exercise 1.4. Find the matrix of the half turn $(x,y) \mapsto (-x,-y)$ and of the reflection $(x,y) \mapsto (y,x)$ across the line $y = x$. Compute $J^2 = JJ$ (the matrix product is defined in the next section, but here it is enough to apply J twice to $\vec{e}_1$ and to $\vec{e}_2$), and check that it is the half turn.

Example 1.5 (covectors). A linear map $\mathbb{R}^n \to \mathbb{R}$ has an $m \times n$ matrix with $m = 1$: a single row. For example, the map

$$\varphi(x,y,z) = 2x - y + 5z$$

is linear, and $\varphi(\vec{e}_1) = 2$, $\varphi(\vec{e}_2) = -1$, $\varphi(\vec{e}_3) = 5$. So its matrix is the row $\begin{pmatrix} 2 & -1 & 5 \end{pmatrix}$, and applying φ to a vector is multiplying the row by the column:

$$\begin{pmatrix} 2 & -1 & 5 \end{pmatrix} \begin{pmatrix} 1 \\ 4 \\ 1 \end{pmatrix} = 2\cdot 1 + (-1)\cdot 4 + 5\cdot 1 = 3,$$

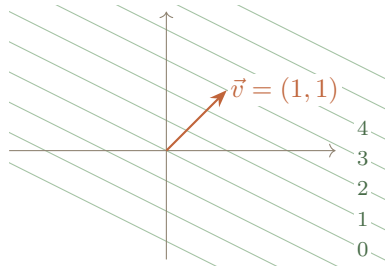
which agrees with $\varphi(1,4,1) = 2 - 4 + 5 = 3$.

Linear maps $\mathbb{R}^n \to \mathbb{R}$ are exactly the *covectors* we have seen before. A covector is a machine that eats a vector and produces a number, linearly, and the row is its recipe: here, "take 2 of the first entry, -1 of the second, and 5 of the third, and add."

Not every function $\mathbb{R}^n \to \mathbb{R}$ is a covector. The length $\vec{v} \mapsto \|\vec{v}\|$ is not linear, since $\|-\vec{v}\| = \|\vec{v}\|$, whereas a linear map would have to send $-\vec{v}$ to the negative

of whatever it sends $\vec{v}$ to. Linearity would force $\|\vec{v}\| = -\|\vec{v}\|$, which fails for every nonzero $\vec{v}$.

There is a good way to picture a covector on $\mathbb{R}^2$. Take the row $\begin{pmatrix}1 & 2\end{pmatrix}$, the covector $(x, y) \mapsto x + 2y$. The vectors it sends to 0 form the line $x + 2y = 0$; the vectors it sends to 1 form the parallel line $x+2y = 1$; and so on. So a covector on $\mathbb{R}^2$ is a family of evenly spaced parallel lines, and its value on a vector $\vec{v}$ counts how many of these lines $\vec{v}$ crosses, with a sign for the direction of crossing.



The vector $\vec{v} = (1, 1)$ starts on the line labeled 0 and ends on the line labeled 3, and indeed $1 + 2 \cdot 1 = 3$.

Exercise 1.6. A covector ψ on $\mathbb{R}^2$ has $\psi(1, 1) = 3$ and $\psi(1, -1) = 1$. Find its row, and compute $\psi(5, 0)$.

1.2 Dot products and matrix products

The point of this section is one formula, which gets used several times later on: for an $n \times n$ matrix A and vectors $\vec{v}, \vec{w}$,

$$(A\vec{v}) \cdot \vec{w} = \vec{v} \cdot (A^T\vec{w}).$$

In words, a matrix on one side of a dot product can hop to the other side, provided it flips over on the way: A^T, the *transpose* of A, is A reflected across its diagonal. The formula is stated precisely as Proposition 1.12 below, after the three facts its proof uses: the dot product is itself a matrix product, matrix products can be regrouped, and transposing a product reverses its order.

Definition 1.7. The *transpose* A^T of an $m \times n$ matrix A is the $n \times m$ matrix with $(A^T)_{ij} = a_{ji}$.

Rows become columns, columns become rows, and the diagonal stays put, smugly. For example,

$$A = \begin{pmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{pmatrix} \qquad \text{has} \qquad A^T = \begin{pmatrix} 1 & 4 \\ 2 & 5 \\ 3 & 6 \end{pmatrix}.$$

The first row of A, $\begin{pmatrix} 1 & 2 & 3 \end{pmatrix}$, became the first column of A^T, and the 2×3 matrix became a 3×2 matrix. In particular, the transpose of a column vector $\vec{v}$ is a row $\vec{v}^T$, for instance $\begin{pmatrix} 2 \\ 1 \end{pmatrix}^T = \begin{pmatrix} 2 & 1 \end{pmatrix}$.

Definition 1.8. Let A be an $m \times n$ matrix and B an $n \times p$ matrix (so the number of columns of A equals the number of rows of B). The *product* AB is the $m \times p$ matrix whose column j is A times column j of B.

The reason for this definition is that it makes $(AB)\vec{v} = A(B\vec{v})$ for every vector $\vec{v}$: multiplying by AB means doing B first and then A. For $\vec{v} = \vec{e}_j$ this is the definition itself, since $(AB)\vec{e}_j$ is column j of AB and $B\vec{e}_j$ is column j of B; for other vectors it follows by linearity. Working out the entries of A times a column gives the entry-by-entry formula

$$(AB)_{ij} = (\text{row } i \text{ of } A) \cdot (\text{column } j \text{ of } B).$$

Example 1.9. Let $A = \begin{pmatrix} 1 & 2 \\ 0 & 3 \end{pmatrix}$ and $B = \begin{pmatrix} 2 & 1 \\ 1 & 0 \end{pmatrix}$. Column by column, $A \begin{pmatrix} 2 \\ 1 \end{pmatrix} = \begin{pmatrix} 4 \\ 3 \end{pmatrix}$ and $A \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$, so

$$AB = \begin{pmatrix} 4 & 1 \\ 3 & 0 \end{pmatrix}, \qquad \text{while} \qquad BA = \begin{pmatrix} 2 & 7 \\ 1 & 2 \end{pmatrix}.$$

The order matters: doing B and then A is usually not the same as doing A and then B. Putting on socks and then shoes is a different experience from the reverse.

So every entry of a matrix product is a dot product. And every dot product is a matrix product:

$$\vec{v} \cdot \vec{w} = \vec{v}^T \vec{w} \qquad (\text{a } 1 \times n \text{ row times an } n \times 1 \text{ column}).$$

This is the same fact we have seen before, that every covector has the form $\vec{h} \mapsto \vec{c} \cdot \vec{h}$ for exactly one vector $\vec{c}$, now in costume: the covector is the row $\vec{c}^T$, and feeding it $\vec{h}$ is the product $\vec{c}^T \vec{h} = \vec{c} \cdot \vec{h}$. Whether we write a covector or a vector, a row or a column, it carries the same information. Which one you write down is a matter of which way you prefer to tilt your head.

Rules for computing with matrices. Matrices of the same size are added, and multiplied by numbers, entry by entry: $(A + B)_{ij} = a_{ij} + b_{ij}$ and $(cA)_{ij} = c\, a_{ij}$. Matrix products and transposes obey a few rules that make computing with them feel almost like computing with numbers. Almost: as the example above showed, AB and BA are usually different, so the order of a product always has to be watched.

Lemma 1.10 (associativity). *If A, B, C are matrices whose sizes allow the products, then $(AB)C = A(BC)$.*

Proof. Compare column j on both sides; write $\vec{c}_j$ for column j of C. By the definition of the product, column j of $(AB)C$ is $(AB)\vec{c}_j$, which is $A(B\vec{c}_j)$ by the property $(AB)\vec{v} = A(B\vec{v})$. Column j of $A(BC)$ is A times column j of BC, which is $A(B\vec{c}_j)$. These agree. □

So the parentheses in a product of several matrices can be dropped, and we may write ABC. (The order still matters; only the grouping does not.)

Lemma 1.11 (transpose rules). *For matrices of sizes that allow the operations, and a number c,*

$$(A + B)^T = A^T + B^T, \qquad (cA)^T = c\, A^T, \qquad \left(A^T\right)^T = A,$$

and

$$(AB)^T = B^T A^T.$$

Proof. Compare entries. The (i, j) entry of $(A + B)^T$ is the (j, i) entry of

$A + B$, which is $a_{ji} + b_{ji}$, and that is also the (i, j) entry of $A^T + B^T$. The second rule is the same argument with $c\,a_{ji}$. For the third, transposing twice swaps the indices twice, which puts every entry back where it started. For the last rule,

$$\begin{aligned}\big((AB)^T\big)_{ij} = (AB)_{ji} &= (\text{row } j \text{ of } A) \cdot (\text{column } i \text{ of } B),\\ \big(B^T A^T\big)_{ij} &= (\text{row } i \text{ of } B^T) \cdot (\text{column } j \text{ of } A^T)\\ &= (\text{column } i \text{ of } B) \cdot (\text{row } j \text{ of } A).\end{aligned}$$

These are the same dot product, taken in the other order. □

The last rule is the one to remember: *transposing reverses the order of a product*. Socks and shoes again: to undo putting on socks and then shoes, take off the shoes first.

For example,

$$\left(\begin{pmatrix}1 & 2\\ 3 & 4\end{pmatrix} + \begin{pmatrix}0 & 5\\ 1 & 0\end{pmatrix}\right)^T = \begin{pmatrix}1 & 7\\ 4 & 4\end{pmatrix}^T = \begin{pmatrix}1 & 4\\ 7 & 4\end{pmatrix} = \begin{pmatrix}1 & 3\\ 2 & 4\end{pmatrix} + \begin{pmatrix}0 & 1\\ 5 & 0\end{pmatrix},$$

and, with A and B from the example above, $(AB)^T = \begin{pmatrix}4 & 3\\ 1 & 0\end{pmatrix} = B^T A^T$, while $A^T B^T = (BA)^T = \begin{pmatrix}2 & 1\\ 7 & 2\end{pmatrix}$ is different.

Now the formula itself.

Proposition 1.12 (moving a matrix across a dot product). *Let A be an $n \times n$ matrix and $\vec{v}, \vec{w} \in \mathbb{R}^n$. Then*

$$(A\vec{v}) \cdot \vec{w} = \vec{v} \cdot (A^T \vec{w}).$$

Proof. Think of $\vec{v}$ and $\vec{w}$ as $n \times 1$ matrices, and of a 1×1 matrix as a number. Writing the dot product as a matrix product, reversing the transpose, and regrouping,

$$(A\vec{v}) \cdot \vec{w} = (A\vec{v})^T \vec{w} = \big(\vec{v}^T A^T\big)\vec{w} = \vec{v}^T \big(A^T \vec{w}\big) = \vec{v} \cdot \big(A^T \vec{w}\big).$$

The second equality is the reversal rule of Lemma 1.11, and the third is associativity (Lemma 1.10). □

This formula is the main reason anyone tolerates the transpose.

Exercise 1.13. Let $A = \begin{pmatrix} 1 & 2 \\ 0 & 3 \end{pmatrix}$, $\vec{v} = (1, 1)$, and $\vec{w} = (2, -1)$. Compute $(A\vec{v}) \cdot \vec{w}$ and $\vec{v} \cdot (A^T\vec{w})$, and check that they agree.

1.3 Symmetric and antisymmetric

Every square matrix is, in exactly one way, a symmetric matrix plus an antisymmetric one (Proposition 1.14 below). For the derivative matrix of a vector field, the antisymmetric part will turn out to be the curl.

A square matrix A is *symmetric* if $A^T = A$ and *antisymmetric* if $A^T = -A$. A symmetric matrix looks the same in a mirror placed along its diagonal. An antisymmetric matrix looks the same except that every entry has changed its mind about its sign. The diagonal entries of an antisymmetric matrix satisfy $a_{ii} = -a_{ii}$, and the only number that is its own negative is 0, so the diagonal is all zeros. The entries below the diagonal then determine those above.

Proposition 1.14. *Every square matrix A can be written in exactly one way as $A = S + K$ with S symmetric and K antisymmetric, namely*

$$S = \tfrac{1}{2}\big(A + A^T\big), \qquad K = \tfrac{1}{2}\big(A - A^T\big).$$

Proof. First, these S and K work. By Lemma 1.11,

$$S^T = \tfrac{1}{2}\big(A^T + (A^T)^T\big) = \tfrac{1}{2}\big(A^T + A\big) = S, \qquad K^T = \tfrac{1}{2}\big(A^T - A\big) = -K,$$

and $S + K = \frac{1}{2}A + \frac{1}{2}A = A$.

Second, no other choice works. Suppose also $A = S' + K'$ with S' symmetric and K' antisymmetric. Then $S - S' = K' - K$; call this matrix M. By Lemma 1.11, $M^T = S^T - S'^T = S - S' = M$, so M is symmetric, and also $M^T = K'^T - K^T = -K' + K = -M$, so M is antisymmetric. Then $M = M^T = -M$, so every entry of M is its own negative, hence 0. So $S' = S$ and $K' = K$. □

K is called the *antisymmetric part* of A, and passing from A to K is called *antisymmetrizing*.

Example 1.15. For the matrix of Example 1.2,

$$\begin{pmatrix} 2 & 0 \\ 1 & 1 \end{pmatrix} = \begin{pmatrix} 2 & \frac{1}{2} \\ \frac{1}{2} & 1 \end{pmatrix} + \begin{pmatrix} 0 & -\frac{1}{2} \\ \frac{1}{2} & 0 \end{pmatrix} = S + \tfrac{1}{2}J.$$

So this slanting map is a symmetric matrix plus half of a quarter turn. Geometrically, a symmetric matrix stretches along perpendicular axes, possibly by negative factors (true, though not proved here), and in the plane an antisymmetric matrix cJ is a quarter turn followed by scaling by c. So every linear map of the plane is a stretch plus a twist. The twist is the part these notes care about.

Exercise 1.16. Let $T\colon \mathbb{R}^2 \to \mathbb{R}^2$ be linear with $T(\vec{e}_1) = (2, 1)$ and $T(\vec{e}_2) = (-1, 3)$. Find its matrix A, compute $T(3, 2)$ two ways (by columns and by rows), and split A into its symmetric and antisymmetric parts.

Exercise 1.17. Show that if K is antisymmetric, then $\vec{v} \cdot (K\vec{v}) = 0$ for every $\vec{v}$. (Use Proposition 1.12.) What does this say about the vector $J\vec{v}$ and the vector $\vec{v}$?

Counting. In three dimensions, and only there, an antisymmetric matrix carries exactly as much information as a vector. Here is why. An antisymmetric $n \times n$ matrix is determined by its entries below the diagonal, of which there are $\frac{n(n-1)}{2}$.

- For $n = 2$ that is *one* number: every antisymmetric 2×2 matrix is

$$\begin{pmatrix} 0 & -c \\ c & 0 \end{pmatrix} = cJ.$$

- For $n = 3$ that is *three* numbers, the same as a vector in $\mathbb{R}^3$. This case gets its own notation, in Definition 1.18 below.

- For $n = 4$ it is *six* numbers: every antisymmetric 4×4 matrix is

$$\begin{pmatrix} 0 & -a & -b & -c \\ a & 0 & -d & -e \\ b & d & 0 & -f \\ c & e & f & 0 \end{pmatrix}$$

  for six numbers a, b, c, d, e, f. Six is neither one nor four, so an antisymmetric 4×4 matrix is not the same thing as a single number or as a single vector in $\mathbb{R}^4$.

Here is the notation for $n = 3$.

Definition 1.18. For a vector $\vec{c} = (c_1, c_2, c_3)$ in $\mathbb{R}^3$, let $[\vec{c}]$ be the antisymmetric matrix

$$[\vec{c}] = \begin{pmatrix} 0 & -c_3 & c_2 \\ c_3 & 0 & -c_1 \\ -c_2 & c_1 & 0 \end{pmatrix}.$$

Every antisymmetric 3×3 matrix K is $[\vec{c}]$ for exactly one vector $\vec{c}$, and $\vec{c}$ can be read off from the entries of K:

$$\vec{c} = \big(k_{32},\ k_{13},\ k_{21}\big). \tag{1.1}$$

That is, c_1 is the entry in row 3, column 2; c_2 is the entry in row 1, column 3; and c_3 is the entry in row 2, column 1. These are exactly the three places where the entries of $[\vec{c}]$ appear with a plus sign in Definition 1.18. For example,

$$[(1, 2, 3)] = \begin{pmatrix} 0 & -3 & 2 \\ 3 & 0 & -1 \\ -2 & 1 & 0 \end{pmatrix}.$$

Three dimensions is the only place where an antisymmetric matrix carries exactly as much information as a vector (since $\frac{n(n-1)}{2} = n$ only for $n = 3$). This coincidence is the entire reason the cross product exists, and a good deal of physics has been built on it without asking whether it was a sensible foundation. Chapter 7 returns to it.

Secretly, a function of two vectors. A square matrix M also defines a function of two vectors, $B(\vec{h}, \vec{k}) = \vec{k} \cdot (M\vec{h})$, which is linear in each input separately. By Proposition 1.12, swapping the inputs transposes the matrix: $B(\vec{k}, \vec{h}) = \vec{k} \cdot (M^T\vec{h})$. So M is antisymmetric exactly when $B(\vec{k}, \vec{h}) = -B(\vec{h}, \vec{k})$ always. Antisymmetric functions of two vectors are what we previously called *2-forms*. The basic one in the plane is $\vec{k} \cdot (J\vec{h}) = h_1k_2 - h_2k_1 = \det(\vec{h}, \vec{k})$, the signed area of the parallelogram on $\vec{h}$ and $\vec{k}$. Nothing below depends on this box, but it is what is really going on.

Chapter 2

The derivative of a function: a row, or a vector

The derivative's merely a row
Of partials, all lined up for show.
Turn it up on its side
And the gradient's your guide
To the uphill, where water won't go.

The derivative of a function $f\colon \mathbb{R}^n \to \mathbb{R}$ at a point $\vec{a}$ is a row of numbers, the partial derivatives of f at $\vec{a}$. Written as a column, the same numbers form the gradient $\nabla f(\vec{a})$.

Recall that $f\colon \mathbb{R}^n \to \mathbb{R}$ is *differentiable* at $\vec{a}$ if there is a covector $df|_{\vec{a}}$ with

$$\lim_{\vec{h}\to\vec{0}} \frac{f(\vec{a}+\vec{h}) - f(\vec{a}) - df|_{\vec{a}}(\vec{h})}{\|\vec{h}\|} = 0.$$

In words: near $\vec{a}$, the change in f is a linear function of the change in position, up to an error that is small even compared with $\|\vec{h}\|$. The derivative is the best impression of f near $\vec{a}$ that a linear function can manage.

As a matrix, $df|_{\vec{a}}$ is a single row, and its entries are the partial derivatives:

$$df|_{\vec{a}} = \begin{pmatrix} \dfrac{\partial f}{\partial x_1}(\vec{a}) & \cdots & \dfrac{\partial f}{\partial x_n}(\vec{a}) \end{pmatrix} = \nabla f(\vec{a})^T.$$

Its partner vector, the column with the same entries, is the gradient, and

$$df|_{\vec{a}}(\vec{h}) = \nabla f(\vec{a}) \cdot \vec{h}.$$

So the derivative of f at $\vec{a}$ can be recorded either as a row (the covector $df|_{\vec{a}}$) or as a column (the vector $\nabla f(\vec{a})$). As $\vec{a}$ moves, we get a *covector field* df and its

partner, the *vector field* ∇f.

What the derivative measures. If f is differentiable at $\vec{a}$, then the chain rule along the line $t \mapsto \vec{a} + t\vec{h}$ gives

$$\frac{d}{dt}\Big|_{t=0} f(\vec{a} + t\vec{h}) = \nabla f(\vec{a}) \cdot \vec{h}. \tag{2.1}$$

So $\nabla f(\vec{a}) \cdot \vec{h}$ is how fast f changes as you pass through $\vec{a}$ with velocity $\vec{h}$. For a unit vector $\vec{u}$ this is the directional derivative $D_{\vec{u}} f(\vec{a})$, and it is largest when $\vec{u}$ points along $\nabla f(\vec{a})$. In other words, the gradient points in the direction in which f goes uphill fastest. Water, which has never been told about gradients, runs the other way.

When does the derivative exist? If the partial derivatives of f are continuous, then f is differentiable (Theorem 2.1 below). Mere existence of the partial derivatives is not enough: partial derivatives can exist at a point where f is not differentiable. Each partial derivative only looks along a single coordinate axis, and a function can behave perfectly along every axis while doing something deplorable along the diagonals. So the test that actually gets used is the following, which we state without proof.

Theorem 2.1 (continuous partials).
If the partial derivatives $\partial f/\partial x_1, \dots, \partial f/\partial x_n$ exist near $\vec{a}$ and are continuous at $\vec{a}$, then f is differentiable at $\vec{a}$.

A function is C^1 on an open set S if its partial derivatives exist and are continuous on S, and C^2 if, in addition, the partial derivatives of its partial derivatives exist and are continuous on S. By Theorem 2.1, a C^1 function is differentiable at every point of S.

Chapter 3

The derivative of a vector field: a matrix

Take a field, its components in tow;
Stack their gradients, row upon row.
That's $D\vec{F}$, and it's square;
If $\vec{F}$'s grad G, I swear,
Then it's symmetric—Hessians are so.

A vector field $\vec{F} = (F_1, \dots, F_n)$ on an open set $S \subseteq \mathbb{R}^n$ is really n ordinary functions $F_1, \dots, F_n$ standing on each other's shoulders in a trench coat. Each F_i has a gradient, and stacking these gradients as rows gives the derivative of $\vec{F}$.

Definition 3.1. Let $\vec{F}$ be C^1 on S, meaning that each F_i is C^1. The *derivative matrix* $D\vec{F}(\vec{a})$ is the $n \times n$ matrix whose i-th row is $\nabla F_i(\vec{a})^T$:

$$D\vec{F}(\vec{a}) = \begin{pmatrix} \dfrac{\partial F_1}{\partial x_1} & \cdots & \dfrac{\partial F_1}{\partial x_n} \\ \vdots & \ddots & \vdots \\ \dfrac{\partial F_n}{\partial x_1} & \cdots & \dfrac{\partial F_n}{\partial x_n} \end{pmatrix}_{\vec{a}}, \qquad \left(D\vec{F}(\vec{a})\right)_{ij} = \frac{\partial F_i}{\partial x_j}(\vec{a}).$$

To remember the layout: the component F_i changes as you go down the rows, and the variable x_j changes as you go across the columns.

Applying (2.1) to each component, which is legitimate because each F_i is differentiable by Theorem 2.1, gives

$$\frac{d}{dt}\Big|_{t=0} \vec{F}(\vec{a} + t\vec{h}) = D\vec{F}(\vec{a})\,\vec{h}. \tag{3.1}$$

So $D\vec{F}(\vec{a})\vec{h}$ is the rate at which the *arrow* $\vec{F}$ changes as you pass through $\vec{a}$ with velocity $\vec{h}$. The matrix $D\vec{F}(\vec{a})$ gives the best linear approximation of $\vec{F}$ near $\vec{a}$: $\vec{F}(\vec{a} + \vec{h}) \approx \vec{F}(\vec{a}) + D\vec{F}(\vec{a})\vec{h}$.

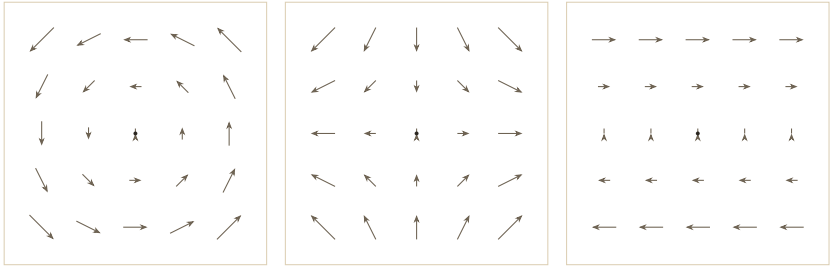
Here are four examples, which together cover most of the personalities a vector field can have.

Example 3.2 (linear fields). If $\vec{F}(\vec{x}) = A\vec{x}$ for a fixed matrix A, then $D\vec{F} = A$ at every point. A linear field is its own best linear approximation, which is either very deep or not deep at all.

Example 3.3 (the swirl). $\vec{F}(x, y) = (-y, x)$ is the linear field with matrix J, so $D\vec{F} = J$ everywhere. The arrow at each point is that point's position vector, turned a quarter turn counterclockwise. Release a cork into this field and it goes round and round forever.

Example 3.4 (a gradient). $\vec{F}(x, y) = (x, -y)$ is ∇G for $G = \frac{1}{2}(x^2 - y^2)$, and $D\vec{F} = \left(\begin{smallmatrix} 1 & 0 \\ 0 & -1 \end{smallmatrix}\right)$. The arrows flow in from above and below and out to the left and right, as if the plane were being pulled apart at the sides. Nothing goes round.

Example 3.5 (the shear). $\vec{F}(x, y) = (y, 0)$: a river flowing to the right, fast near the top, slow near the bottom, and backwards below the x-axis. Here $D\vec{F} = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$. Every arrow is horizontal. Nothing appears to go round. Hold on to that thought; it will turn out to be misleading.



swirl $(-y, x)$ gradient $(x, -y)$ shear $(y, 0)$

3.1 Is it a gradient? A first look

If $\vec{F}$ is a gradient, then its derivative matrix $D\vec{F}$ is symmetric. This already shows that some of the fields above are not gradients. Here is the reason. Suppose $\vec{F} = \nabla G$. Then row i of $D\vec{F}$ is the gradient of $\partial G/\partial x_i$, and

$$(D\vec{F})_{ij} = \frac{\partial}{\partial x_j}\frac{\partial G}{\partial x_i}.$$

This is the matrix of second partial derivatives of G, the *Hessian*. If mixed partials can be taken in either order (they can, for C^2 functions; Chapter 5), the Hessian is symmetric. So:

if $\vec{F}$ is a gradient, then $D\vec{F}$ is symmetric.

Now check this against the examples. The gradient $(x, -y)$ has a symmetric derivative, as it must. The swirl has $D\vec{F} = J$, which is as far from symmetric as a matrix can get: it is antisymmetric. So the swirl is not a gradient, and we did not have to integrate anything to find out. The shear has $D\vec{F} = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$, which is not symmetric. So the shear, despite having no visible whirlpool, is not a gradient either.

The symmetric part of $D\vec{F}$ tells us nothing here, since every matrix has one. It is a nonzero antisymmetric part that is evidence of a crime: it proves that $\vec{F}$ is not a gradient. So the antisymmetric part gets a name.

Exercise 3.6. Let $\vec{F}(x, y) = (x^2 y,\ x + y^3)$. Compute $D\vec{F}$ at the points $(1, 2)$ and $(2, 1)$. At which of these points is $D\vec{F}$ symmetric? What does this say about whether $\vec{F}$ is a gradient?

Chapter 4

Curl: the antisymmetric part

Take $D\vec{F}$ and subtract its transpose;
What is left is the curl, and it shows
How a paddle wheel twirls,
Even where nothing swirls,
As the shear on a lazy stream goes.

Definition 4.1. Let $\vec{F}$ be C^1 on an open set $S \subseteq \mathbb{R}^n$. The *curl* of $\vec{F}$ at $\vec{a}$ is the antisymmetric matrix

$$\operatorname{curl}\vec{F}(\vec{a}) = D\vec{F}(\vec{a}) - \big(D\vec{F}(\vec{a})\big)^T, \qquad \big(\operatorname{curl}\vec{F}\big)_{ij} = \frac{\partial F_i}{\partial x_j} - \frac{\partial F_j}{\partial x_i}.$$

This is twice the antisymmetric part of $D\vec{F}$. The factor of 2 is there to avoid a $\frac{1}{2}$ that would otherwise follow the curl around for the rest of its life, appearing in every formula like a relative who was never formally invited. With the factor of 2, the formulas below match the ones found in every other book.

In the plane. For $\vec{F} = (P, Q)$, the curl is an antisymmetric 2×2 matrix, so it is one number times J:

$$\begin{aligned}\operatorname{curl}\vec{F} &= \begin{pmatrix} P_x & P_y \\ Q_x & Q_y \end{pmatrix} - \begin{pmatrix} P_x & Q_x \\ P_y & Q_y \end{pmatrix} \\ &= \begin{pmatrix} 0 & P_y - Q_x \\ Q_x - P_y & 0 \end{pmatrix} = (Q_x - P_y)\,J.\end{aligned}$$

The number $Q_x - P_y$ is called the *scalar curl* of $\vec{F}$. In the plane, this number is what people usually mean when they say "the curl." It is zero exactly when $\operatorname{curl}\vec{F}$ is the zero matrix.

Example 4.2. (a) Swirl, $(-y, x)$: $Q_x - P_y = 1 - (-1) = 2$.

(b) Gradient, $(x, -y)$: $Q_x - P_y = 0 - 0 = 0$.

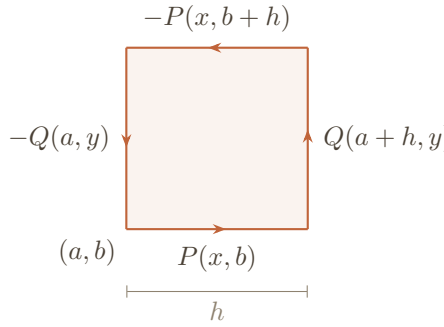
(c) Shear, $(y, 0)$: $Q_x - P_y = 0 - 1 = -1$.

(d) Linear, $A\vec{x}$: curl $\vec{F} = A - A^T$, which is the same matrix at every point, and is zero exactly when A is symmetric.

Exercise 4.3. Compute the scalar curl $Q_x - P_y$ of $\vec{F} = (x^2 y,\ x + y^3)$ from Exercise 3.6. Where in the plane is it zero?

4.1 What curl measures

In the plane, the scalar curl is *circulation per unit area* (Proposition 4.4 below). The definition of the curl is an algebraic statement about a matrix; this is what it means geometrically. Take a tiny square, go around its edge counterclockwise, and add up the work $\vec{F}$ does along the way. That number is the *circulation* of $\vec{F}$ around the square. Now divide it by the area of the square, and let the square shrink.



Proposition 4.4 (circulation per unit area). *Let $\vec{F} = (P, Q)$ be C^1 near (a, b), and let C_h be the boundary of the square $[a, a+h] \times [b, b+h]$, traversed counterclockwise. Then*

$$\lim_{h \to 0^+} \frac{1}{h^2} \int_{C_h} \vec{F} \cdot d\vec{r} = Q_x(a, b) - P_y(a, b).$$

Here is the idea of the proof, before the ε's. On the bottom edge, moving right, you collect about $h\,P$. On the top edge, moving left, you give back about $h\,P$,

but with P evaluated a distance h higher up. The difference is about $-h \cdot h\, P_y$. In the same way, the two vertical sides together contribute about $+h \cdot h\, Q_x$.

Proof. Parametrize the four sides as segments. The bottom and top combine, as do the right and left, to give

$$\int_{C_h} \vec{F} \cdot d\vec{r} = \int_a^{a+h} \big[P(x,b) - P(x,b+h)\big]\,dx + \int_b^{b+h} \big[Q(a+h,y) - Q(a,y)\big]\,dy.$$

Let $\varepsilon > 0$. Since P_y and Q_x are continuous at (a,b), for all small enough $h > 0$ every point $\vec{x}$ of the square satisfies $|P_y(\vec{x}) - P_y(a,b)| < \varepsilon$ and $|Q_x(\vec{x}) - Q_x(a,b)| < \varepsilon$. For each x, the mean value theorem in the y variable gives $P(x,b+h) - P(x,b) = h\, P_y(x,\eta)$ for some η between b and $b+h$, so

$$\big|P(x,b) - P(x,b+h) + h\, P_y(a,b)\big| < h\varepsilon.$$

Integrating over an interval of length h shows that the first integral is within $h^2\varepsilon$ of the number $-h^2 P_y(a,b)$. In the same way the second is within $h^2\varepsilon$ of $h^2 Q_x(a,b)$. Dividing by h^2, the quotient is within 2ε of $Q_x(a,b) - P_y(a,b)$. □

So the scalar curl is circulation per unit area. Here is a physical way to say it. Drop a tiny paddle wheel into the field at (a,b), pinned so that it can spin but not drift. If the curl at (a,b) is positive, the wheel spins counterclockwise; if it is negative, the wheel spins clockwise; and if it is zero, the wheel does not spin at all.

Exercise 4.5. Compute the circulation of the shear $(y,0)$ around the boundary of $[0,h] \times [0,h]$, counterclockwise, directly from the definition of the line integral. Compare with Proposition 4.4.

The shear, revisited. The shear $(y,0)$ has curl -1, so the paddle wheel turns clockwise. This is not because anything in the river goes round. It is because the water across the top of the wheel moves right faster than the water across the bottom, and the wheel has no choice in the matter. Rivers are like this: there

are no whirlpools anywhere, and yet every leaf on the surface slowly rotates. So having nonzero curl is not the same as having arrows that go in circles. Curl measures the local, infinitesimal tendency to twist, which is a subtler thing, and it is the right thing for deciding whether $\vec{F}$ is a gradient.

It also works the other way: a field can have arrows that go in circles and still have curl zero. The angle field of Chapter 6 does exactly that, and it causes all the trouble there is to cause.

Chapter 5

Gradients have zero curl

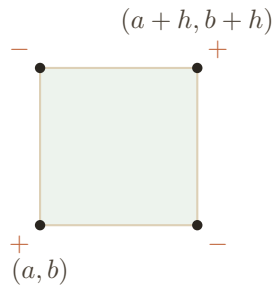
When a field is the grad of some G,
Then its curl is as zero can be,
Since the partials, when mixed,
Never care which came first:
Group the corners two ways: they agree.

Chapter 3 showed that the statement "if $\vec{F}$ is a gradient, then $D\vec{F}$ is symmetric" comes down to one fact about second partial derivatives. Here is that fact. We write f_{xy} for $(f_x)_y$, meaning that we differentiate first in x and then in y.

Theorem 5.1 (mixed partials). *Let f be a function of (x, y) such that f_x, f_y, f_{xy}, and f_{yx} exist near (a, b), and f_{xy} and f_{yx} are continuous at (a, b). Then $f_{xy}(a, b) = f_{yx}(a, b)$.*

The idea is to compute one number in two ways. Take a small square with a corner at (a, b) and form the alternating sum of f at its four corners:

$$\Delta(h) = f(a + h, b + h) - f(a + h, b) - f(a, b + h) + f(a, b).$$



Group the four terms as (top minus bottom) on the right, minus (top minus bottom) on the left, and the mean value theorem turns $\Delta(h)$ into $h^2 f_{xy}$ at some point of the square. Group them as (right minus left) on top, minus (right minus

left) on the bottom, and it turns $\Delta(h)$ into $h^2 f_{yx}$ at some other point. The four numbers do not care how they are grouped. That is the whole proof; the rest is making the mean value theorem do the paperwork.

Proof. Let $h \neq 0$ be small.

First grouping. Let $\varphi(x) = f(x, b+h) - f(x, b)$, so that $\Delta(h) = \varphi(a+h) - \varphi(a)$. Since f_x exists near (a, b), the mean value theorem gives ξ between a and $a + h$ with

$$\Delta(h) = h\,\varphi'(\xi) = h\big[f_x(\xi, b+h) - f_x(\xi, b)\big].$$

Since f_{xy} exists near (a, b), the mean value theorem applied to $y \mapsto f_x(\xi, y)$ gives η between b and $b + h$ with $\Delta(h) = h^2 f_{xy}(\xi, \eta)$.

Second grouping. Let $\psi(y) = f(a + h, y) - f(a, y)$, so that $\Delta(h) = \psi(b + h) - \psi(b)$. The same two steps in the other order, using f_y and f_{yx}, give $\Delta(h) = h^2 f_{yx}(\xi', \eta')$ for a point (ξ', η') of the square.

So $f_{xy}(\xi, \eta) = f_{yx}(\xi', \eta')$. As $h \to 0$, both points tend to (a, b), and continuity of f_{xy} and f_{yx} at (a, b) gives $f_{xy}(a, b) = f_{yx}(a, b)$. □

Remark 5.2. The continuity hypothesis is not decoration. The function

$$f(x, y) = \frac{xy(x^2 - y^2)}{x^2 + y^2}, \qquad f(0, 0) = 0,$$

has $f_x(0, y) = -y$ and $f_y(x, 0) = x$, so $f_{xy}(0, 0) = -1$ and $f_{yx}(0, 0) = 1$. The theorem survives this because its second partials are not continuous at the origin. The function was constructed for no other purpose than to make this point, and has been making it in textbooks for well over a century.

Exercise 5.3. For $f(x, y) = x^3 y^2 + \sin(xy)$, compute f_{xy} and f_{yx}, and check that they agree.

In $\mathbb{R}^n$, apply Theorem 5.1 to $(s, t) \mapsto G(\vec{a} + s\vec{e}_i + t\vec{e}_j)$ for $i \neq j$. If G is C^2, this gives $\dfrac{\partial}{\partial x_j}\dfrac{\partial G}{\partial x_i} = \dfrac{\partial}{\partial x_i}\dfrac{\partial G}{\partial x_j}$: the Hessian is symmetric.

Corollary 5.4 (gradients have zero curl). *If G is C^2 on an open set S, then*

$\operatorname{curl}(\nabla G) = 0$ *on* S. *Equivalently, if* $\vec{F} = \nabla G$ *with* $\vec{F}$ C^1, *then* $D\vec{F}$ *is symmetric:* $\partial F_i/\partial x_j = \partial F_j/\partial x_i$ *for all* i, j. *In the plane,* $P_y = Q_x$.

Proof. $D(\nabla G)$ is the Hessian of G, which is symmetric, so its antisymmetric part is zero. For the second form: if $\vec{F} = \nabla G$ is C^1, then the partial derivatives of G are C^1, so G is C^2. □

This gives a test with teeth. To show that $\vec{F}$ is *not* a gradient, it is enough to compute one entry of $\operatorname{curl} \vec{F}$ and find that it is not zero. For instance, the swirl ($Q_x - P_y = 2$) and the shear ($Q_x - P_y = -1$) are not gradients. This needed no loops, no integrals, and no couch.

> **The slogan.** Taking a function to its gradient, and taking a vector field to its curl, both follow the recipe "take the derivative, and keep the interesting part." Corollary 5.4 says that doing it twice gives zero. In the language of differential forms, which we have glimpsed before, this is written $d \circ d = 0$, and it is one of the few equations in mathematics that is both very short and very important.

Definition 5.5. A C^1 vector field $\vec{F}$ on S is *closed* if $\operatorname{curl} \vec{F} = 0$ on S, and *exact* (or *conservative*, or a *gradient field*) if $\vec{F} = \nabla G$ for some function G on S.

In this language, Corollary 5.4 says that every exact field is closed. The obvious question is whether every closed field is exact. The answer is a firm "it depends where you are standing."

Chapter 6

Closed fields on a disk are gradients

On a disk where the curl is all naught,
A potential is readily got:
Integrate along rays
From the center. Malaise
Only comes when there's a hole in the plot.

The main result of this chapter is the converse of Corollary 5.4, on a disk: if curl $\vec{F} = 0$ on an open disk (or ball, or more generally a star-shaped set), then $\vec{F} = \nabla G$ for an explicit function G (Theorem 6.3). Before proving it, an example shows that some hypothesis on the region is needed: on a plane with a hole, it fails.

6.1 The angle field, or: the hole is the problem

Example 6.1 (the angle field). On the punctured plane $\mathbb{R}^2 \setminus \{\vec{0}\}$, let

$$\vec{F} = (P, Q) = \left(\frac{-y}{x^2+y^2}, \frac{x}{x^2+y^2}\right).$$

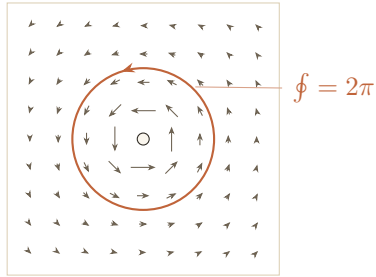
This is the swirl, divided by the squared distance to the origin: the arrows still go round, but they get weaker as you move out. A computation gives

$$P_y = \frac{y^2 - x^2}{(x^2+y^2)^2} = Q_x,$$

so $\vec{F}$ is closed. But along the unit circle $\vec{r}(t) = (\cos t, \sin t)$, $0 \le t \le 2\pi$, $\vec{F}(\vec{r}(t)) \cdot \vec{r}\,'(t) = \sin^2 t + \cos^2 t = 1$, so

$$\oint \vec{F} \cdot d\vec{r} = 2\pi \neq 0.$$

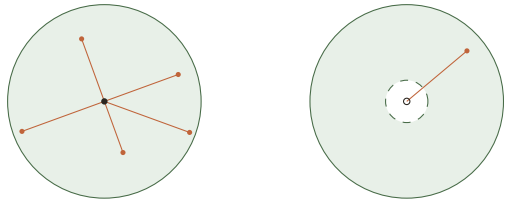
The integral of a gradient around a closed curve is zero, so $\vec{F}$ is not a gradient on $\mathbb{R}^2 \setminus \{\vec{0}\}$.



So the arrows go round, the curl is zero at every point where the field is defined, and still the field is not a gradient. Every tiny paddle wheel sits perfectly still, and yet a trip once around the origin collects 2π worth of work. The explanation is the hole at the origin. The circle cannot be shrunk to a point without passing through the origin, where $\vec{F}$ is not defined, and Proposition 4.4 only knows about tiny squares that lie where $\vec{F}$ lives. The curl gathers local information, and the hole is a global fact that no local measurement can see.

If there is no hole, the problem goes away: on a disk, every closed field is a gradient. The theorem below proves this for a slightly larger class of regions, because the proof works there with no extra effort.

Definition 6.2. An open set $S \subseteq \mathbb{R}^n$ is *star-shaped about* $\vec{0}$ if for every $\vec{x} \in S$, the whole segment $\{t\vec{x} : 0 \leq t \leq 1\}$ lies in S. Every open disk or ball centered at $\vec{0}$ is star-shaped about $\vec{0}$. The punctured plane is not, since every segment from $\vec{0}$ begins at the missing point.



a disk: star-shaped

with a hole: not

6.2 The Poincaré lemma

Theorem 6.3 (Poincaré lemma, first version). *Let $S \subseteq \mathbb{R}^n$ be open and star-shaped about $\vec{0}$, and let $\vec{F}$ be a C^1 vector field on S with* $\operatorname{curl} \vec{F} = 0$. *Then $\vec{F} = \nabla G$, where*

$$G(\vec{x}) = \int_0^1 \vec{F}(t\vec{x}) \cdot \vec{x}\, dt.$$

Before the proof, here is where the formula comes from. The curve $\vec{r}(t) = t\vec{x}$, $0 \le t \le 1$, runs along the segment from $\vec{0}$ to $\vec{x}$, with velocity $\vec{r}\,'(t) = \vec{x}$. So $G(\vec{x})$ is the line integral of $\vec{F}$ along that segment. If $\vec{F}$ is going to be ∇G for some G with $G(\vec{0}) = 0$, the fundamental theorem for line integrals *forces* this formula. So the formula for G is not a clever guess; it is the only possible guess. The work in the proof is checking that the guess actually works, and that check is where the hypothesis $\operatorname{curl} \vec{F} = 0$ gets used, exactly once.

The proof needs one fact that these notes do not prove: under the hypotheses of the lemma below, a derivative can be moved inside an integral.

Lemma 6.4 (differentiating under the integral). *Let $S \subseteq \mathbb{R}^n$ be open, and let $g(\vec{x}, t)$ be defined for $\vec{x} \in S$ and $t \in [0, 1]$, such that g and $\partial g/\partial x_i$ are continuous. Then*

$$\frac{\partial}{\partial x_i} \int_0^1 g(\vec{x}, t)\, dt = \int_0^1 \frac{\partial g}{\partial x_i}(\vec{x}, t)\, dt.$$

Proof. Write $g(\vec{x}, t) = \vec{F}(t\vec{x}) \cdot \vec{x} = \sum_j F_j(t\vec{x})\, x_j$, and fix i. For fixed t, the partial derivative of $\vec{x} \mapsto F_j(t\vec{x})$ with respect to x_i is $t\, \dfrac{\partial F_j}{\partial x_i}(t\vec{x})$, by the chain rule along the line in the x_i direction. By the product rule,

$$\frac{\partial g}{\partial x_i}(\vec{x}, t) = F_i(t\vec{x}) + t \sum_j x_j\, \frac{\partial F_j}{\partial x_i}(t\vec{x}),$$

which is continuous in $(\vec{x}, t)$. Now use that $\vec{F}$ is closed: $\partial F_j/\partial x_i = \partial F_i/\partial x_j$. So

$$\sum_j x_j\, \frac{\partial F_j}{\partial x_i}(t\vec{x}) = \sum_j x_j\, \frac{\partial F_i}{\partial x_j}(t\vec{x}) = \nabla F_i(t\vec{x}) \cdot \vec{x} = \frac{d}{dt} F_i(t\vec{x}),$$

the last step by the chain rule along $t \mapsto t\vec{x}$. Hence

$$\frac{\partial g}{\partial x_i}(\vec{x}, t) = F_i(t\vec{x}) + t\,\frac{d}{dt}F_i(t\vec{x}) = \frac{d}{dt}\big[t\,F_i(t\vec{x})\big].$$

By Lemma 6.4 and the fundamental theorem of calculus,

$$\frac{\partial G}{\partial x_i}(\vec{x}) = \int_0^1 \frac{d}{dt}\big[t\,F_i(t\vec{x})\big]\,dt = 1\cdot F_i(\vec{x}) - 0\cdot F_i(\vec{0}) = F_i(\vec{x}).$$

So the partial derivatives of G exist and are the continuous functions F_i. By Theorem 2.1, G is differentiable, with $\nabla G = \vec{F}$. □

Example 6.5. $\vec{F} = (2xy, x^2)$ has $P_y = 2x = Q_x$, so it is closed on all of $\mathbb{R}^2$, which is star-shaped. The theorem produces

$$G(x,y) = \int_0^1 \big(2t^2xy\cdot x + t^2x^2\cdot y\big)\,dt = \int_0^1 3t^2x^2y\,dt = x^2y,$$

and indeed $\nabla(x^2y) = (2xy, x^2)$.

Example 6.6. For a linear field $\vec{F}(\vec{x}) = A\vec{x}$ with A symmetric, the theorem gives $G(\vec{x}) = \int_0^1 (tA\vec{x})\cdot\vec{x}\,dt = \frac{1}{2}\,\vec{x}\cdot(A\vec{x})$. Check: $\nabla G = \frac{1}{2}(A + A^T)\vec{x} = A\vec{x}$, using Proposition 1.12. For $A = \left(\begin{smallmatrix}1 & 0\\ 0 & -1\end{smallmatrix}\right)$ this recovers $G = \frac{1}{2}(x^2 - y^2)$ from Example 3.4.

Remark 6.7. (a) The same proof works for a set that is star-shaped about any point $\vec{p}$, using the segments from $\vec{p}$.

(b) Putting the pieces together: on an open ball, for a C^1 field $\vec{F}$,

$$\vec{F} \text{ is a gradient} \iff \operatorname{curl}\vec{F} = 0 \iff \frac{\partial F_i}{\partial x_j} = \frac{\partial F_j}{\partial x_i} \text{ for all } i, j.$$

The angle field shows that "ball" cannot be replaced by "punctured plane."

(c) Where the angle field *is* defined on a star-shaped set, say the half-plane $\{x > 0\}$, it is a gradient there (Exercise 6.10). The trouble was never the field itself. It was the hole.

Exercise 6.8. For $\vec{F} = (y^2,\ 2xy + 3y^2)$, compute $D\vec{F}$ and decide whether $\vec{F}$ is a gradient on $\mathbb{R}^2$. If it is, find G using the formula in Theorem 6.3, and check

the answer.

Exercise 6.9. For $\vec{F} = (yz,\, xz,\, xy)$ on $\mathbb{R}^3$, compute $D\vec{F}$, show that it is symmetric, and find a function G with $\nabla G = \vec{F}$.

Exercise 6.10. Check the computation $P_y = Q_x$ for the angle field. Why does Theorem 6.3 not apply to it on $\mathbb{R}^2 \setminus \{\vec{0}\}$? Does it apply on the half-plane $\{x > 0\}$?

Chapter 7

In three dimensions: the curl vector

In three dimensions, by chance,
A twist and a vector can dance,
Each three numbers in size;
So the curl, in disguise,
And the cross product join the romance.

Chapter 1 noted that $\mathbb{R}^3$ is the one place where an antisymmetric matrix is the same amount of information as a vector. Here is what that coincidence buys. For $\vec{F} = (P, Q, R)$, the curl is the antisymmetric matrix

$$\operatorname{curl} \vec{F} = \begin{pmatrix} 0 & P_y - Q_x & P_z - R_x \\ Q_x - P_y & 0 & Q_z - R_y \\ R_x - P_z & R_y - Q_z & 0 \end{pmatrix}.$$

By Definition 1.18, it is $[\vec{c}\,]$ for exactly one vector $\vec{c}$, and reading off the entries as in (1.1) (row 3 column 2, then row 1 column 3, then row 2 column 1) gives

$$\operatorname{curl} \vec{F} = [\vec{c}\,], \qquad \vec{c} = \big(R_y - Q_z,\ P_z - R_x,\ Q_x - P_y\big).$$

This $\vec{c}$ is the *curl vector* of $\vec{F}$, often written $\nabla \times \vec{F}$, and in most books it is simply called the curl. Its third entry is the scalar curl $Q_x - P_y$ from the plane, which measures the twisting of $\vec{F}$ in the xy-plane. In the same way, the first and second entries measure the twisting in the yz-plane and the zx-plane. The curl vector points along the axis about which a tiny paddle wheel would most like to spin.

The cross product. For vectors $\vec{v}, \vec{w} \in \mathbb{R}^3$, the *cross product* is

$$\vec{v} \times \vec{w} = [\vec{v}\,]\,\vec{w} = \big(v_2 w_3 - v_3 w_2,\ v_3 w_1 - v_1 w_3,\ v_1 w_2 - v_2 w_1\big).$$

In words: turn $\vec{v}$ into its antisymmetric matrix $[\vec{v}]$, and apply that matrix to $\vec{w}$. Two familiar properties of the cross product follow from antisymmetry.

- $\vec{v} \times \vec{w}$ is perpendicular to $\vec{w}$, because $\vec{w} \cdot (K\vec{w}) = 0$ for every antisymmetric K (Exercise 1.17).
- $\vec{v} \times \vec{w}$ is perpendicular to $\vec{v}$, because by Proposition 1.12, $\vec{v} \cdot ([\vec{v}]\vec{w}) = ([\vec{v}]^T\vec{v}) \cdot \vec{w} = -([\vec{v}]\vec{v}) \cdot \vec{w} = 0$; multiplying out shows $[\vec{v}]\vec{v} = \vec{0}$.

In this notation, applying the curl matrix to a vector $\vec{u}$ gives $[\vec{c}]\vec{u} = \vec{c} \times \vec{u}$, the cross product of the curl vector with $\vec{u}$.

Example 7.1. The three-dimensional swirl $\vec{F} = (-y, x, 0)$ has curl vector $(0, 0, 2)$, pointing straight up the axis of rotation. And indeed $\vec{e}_3 \times (x, y, z) = (-y, x, 0)$, so the swirl is the field $\vec{x} \mapsto \vec{e}_3 \times \vec{x}$: take the cross product with the axis. Its curl vector, $2\vec{e}_3$, is twice the axis. The factor of 2 is the same one that appeared in Definition 4.1, and it is not going anywhere.

In $\mathbb{R}^n$ for $n \neq 3$, the count $\frac{n(n-1)}{2}$ is not n, and there is no curl vector and no cross product. The antisymmetric matrix is the notion that works in every dimension, and the cross product is what it looks like in the one dimension where it can pass for a vector.

Exercise 7.2. Compute $(1, 2, 3) \times (0, 1, 0)$ as $[\vec{v}]\vec{w}$, and check that the answer is perpendicular to both vectors. Then find the curl vector of $\vec{F} = (0,\ 0,\ xy)$.