Ideas from Vector Calculus

Kurt Bryan

Most of the facts I state below are for functions of two or three variables, but with noted exceptions all are true for functions of n variables.

0.1 Tangent Line Approximation

If $f(x)$ is a C^2 function (that is, twice differentiable with continuous second derivative) then

$$f(x) = f(x_0) + f'(x_0)(x - x_0) + O(|x - x_0|^2),$$

or, if we let $df = f(x) - f(x_0)$ and $dx = x - x_0$,

$$df = f'(x_0)dx + O(dx^2).$$

For a function of two variables, $f(x, y)$,

$$f(x, y) = f(x_0, y_0)+f_x(x_0, y_0)(x-x_0)+f_y(x_0, y_0)(y-y_0)+O(|x-x_0|^2+|y-y_0|^2)$$

or $df = f_x dx + f_y dy + O(dx^2 + dy^2)$.

0.2 Parametric Curves and Surfaces

To specify a curve parametrically means to specify each coordinate variable as a function of a single independent parameter, e.g., $x = t + 2$, $y = \sin(t)$. This can also be written in a more compact vector-valued function notation, $\mathbf{r}(t) = (t+2, \sin(t))$. Of course the idea extends to three or more dimensions.

If $\mathbf{r}(t)$ is a curve then $\mathbf{r}'(t)$ is the vector

$$\lim_{dt\to 0} \frac{\mathbf{r}(t + dt) - \mathbf{r}(t)}{dt}$$

and is computed in the obvious manner—differentiate component by component. It's not hard to see that $\mathbf{r}'(t)$ is tangent to the curve $\mathbf{r}(t)$—just visualize the points $\mathbf{r}(t)$ and $\mathbf{r}(t + dt)$ on the curve, and let dt approach zero. The vector $\mathbf{r}(t + dt) - \mathbf{r}(t)$ is a secant vector joining the points $\mathbf{r}(t)$ and $\mathbf{r}(t + dt)$; dividing by dt just rescales it. In the limit that $dt \to 0$ it's intuitively clear that the limit defines a tangent vector.

In three (or more) dimensions we can specify two-dimensional surfaces parametrically, by giving x, y, and z as functions of *two* other variables, say s and t. For example, if $0 \leq s \leq \pi$ and $0 \leq t \leq 2\pi$ then

$$\begin{aligned} x &= \sin(s)\cos(t), \\ y &= \sin(s)\sin(t), \\ z &= \cos(s) \end{aligned}$$

specifies the unit sphere in three dimensions.

If a surface is given parametrically as a vector-valued function $\mathbf{r}(s,t)$, then the vectors $\mathbf{r}_s(s,t)$ and $\mathbf{r}_t(s,t)$ are vectors which are tangent to the surface $\mathbf{r}$ at the specified point. You can see this in exactly the same way that you see that $\mathbf{r}'(t)$ is a tangent vector to the one-dimensional curve $\mathbf{r}(t)$.

0.3 The Gradient and Directional Derivatives

The gradient of $f(x,y)$ is a vector and is written ∇f. The definition is

$$\nabla f(x,y) = (f_x(x,y), f_y(x,y))$$

or, if you like the $\mathbf{i}$ and $\mathbf{j}$ notation, $\nabla f(x,y) = f_x(x,y)\mathbf{i} + f_y(x,y)\mathbf{j}$.

IMPORTANT: Geometrically, $\nabla f(x_0, y_0)$ is a vector that points in the direction of steepest increase for the function f at the point (x_0, y_0).

If $\mathbf{n}$ is a unit vector then $\nabla f(x_0,y_0)\cdot\mathbf{n}$ is the rate of increase of the function f in the direction $\mathbf{n}$ at the point (x_0, y_0); this is called the *directional derivative* of f in the direction $\mathbf{n}$ and is usually written as $\frac{\partial f}{\partial \mathbf{n}}$. In two dimensional xy space, if $\mathbf{n} = (1,0)$ or $(0,1)$, respectively, then the directional derivative is just the partial derivative with respect to the corresponding variable, e.g., if $\mathbf{n} = (1,0)$ then $\frac{\partial f}{\partial \mathbf{n}} = \frac{\partial f}{\partial x}$. Similar results are true in any other dimension.

Example: Suppose $f(x,y) = x^2 - xy$ and $\mathbf{n} = (\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}})$. What is $\frac{\partial f}{\partial \mathbf{n}}(1,4)$? First, compute the gradient of f,

$$\nabla f(x,y) = (2x - y, -x).$$

At the point $(1,4)$ $\nabla f(1,4) = (-2,-1)$, so the direction derivative is $(-2)(\frac{1}{\sqrt{2}}) + (-1)(-\frac{1}{\sqrt{2}}) = \frac{-1}{\sqrt{2}}$.

0.4 Vector Fields

A vector field is a function that takes any point (x, y) in the plane and assigns to it a vector. For example, $\mathbf{v}(x, y) = y\mathbf{i} - x\mathbf{j}$ is an example of a vector field. On important example of a vector field is the gradient of a function. If $f(x, y)$ is a function then $\nabla f(x, y)$ is a vector field. You can plot a vector field by choosing a large number of points in the plane, computing the vector associated with each, and then plotting the vectors, each with its tail at the appropriate point:

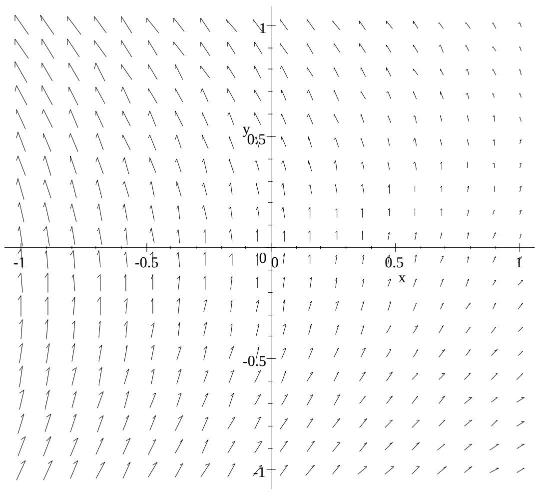

Vector fields are used to describe the velocity or more generally the "flow" of some fluid or other quantity in space, electric and magnetic fields, gravitational fields, and generally any quantity which is a vector and varies from point to point in space.

0.5 Line Integrals

Suppose you have a rope which is 5 meters long and has a linear density of 0.5 kg per meter. What is the mass of the rope? The answer is obvious, 5 meters times 0.5 kg per meter gives 2.5 kg. If density is constant then mass is density times length.

Now suppose that the rope has a variable density—how do we compute the mass? The idea is simple: chop the rope up into many small pieces; each piece has an almost constant density (at least if density varies continuously along the rope), so find the mass of each little piece as above (density times length) and add up the masses. As an example, suppose that the rope is described parametrically in three dimensions by $\mathbf{r}(t) = (t, t^2 + t, \sin(t))$ for $0 \le t \le 2$. Suppose that the density along the rope is $\rho(t) = t + 1$. If we change t by a little bit dt then the resulting change in $\mathbf{r}$ is

$$d\mathbf{r} = \mathbf{r}'(t)\,dt = (1, 2t + 1, \cos(t))\,dt,$$

an infinitesimal vector. Its length ds is

$$ds = |d\mathbf{r}| = \sqrt{2 + 4t + 4t^2 + \cos^2(t)}\,dt.$$

On such a short piece of the rope the density is virtually constant. Therefore the mass of this short piece should be nearly $\rho(t)\,ds$. The total mass is obtained by adding up the mass of each piece,

$$\int_0^2 \rho(t)\,ds = \int_0^2 (t+1)\sqrt{2 + 4t + 4t^2 + \cos^2(t)}\,dt,$$

whatever that works out to be.

More abstractly, we can talk about integrating a function f over a curve C which means computing the integral

$$\int_C f\,ds.$$

In order to actually compute the integral we'd have to parameterize C with some function $\mathbf{r}(t)$ with appropriate limits and then set $ds = |\mathbf{r}'(t)|\,dt$. The parameterization of C that we choose doesn't change the value of the integral (as long as $r' \neq 0$ at any point).

One common use of line integrals is for computing work. Suppose a particle follows a path $\mathbf{r}(t)$ through a vector field F which exerts a force on the particle. Recall also that work is the dot product of the displacement vector with the force, *if the force is constant.* In this case, as the particle moves from t to $t + dt$ the displacement is $\mathbf{r}'(t)\,dt$, and so the work done over this short time interval is $\mathbf{r}'(t) \cdot \mathbf{F}\,dt$. The total work from time t_1 to t_2 is

$$\int_{t_1}^{t_2} \mathbf{r}'(t) \cdot \mathbf{F}(\mathbf{r}(t))\,dt.$$

0.6 Flux over a Curve

For this concept we restrict our attention to two-dimensional space. Later we'll generalize to three dimensions. Suppose that $\mathbf{v}(x, y)$ is a vector field that describes the velocity of some fluid flowing around in two-dimensional space; we'll assume that $\mathbf{v}$ does not depend on time, so the flow is "steady-state", and that the corresponding fluid is incompressible. Note the $\mathbf{v}$ has units of length per time. The question of interest is this: Given a curve C in the plane, how much fluid per unit time is crossing the curve C? The answer should be in square units of fluid per time.

We'll start with a simple case in which $\mathbf{v}$ is constant and C is a line segment, say $\mathbf{v} = (2, 3)$ meters per second. Let S be a line segment with unit normal vector $\mathbf{n}$ as illustrated below:

I'll use $|S|$ to denote the length of S. The shaded parallelogram in the figure represents the fluid that crossed S from time t to time $t + 1$. The area of the parallelogram is, from high school geometry, $|S||\mathbf{v}|\sin(\frac{\pi}{2} - \theta)$ or equivalently, $|S||\mathbf{v}|\cos(\theta)$, which is just $|S|\mathbf{v}\cdot\mathbf{n}$ (here $|\mathbf{v}|$ is the magnitude of $\mathbf{v}$, as a vector). This is the rate at which fluid is crossing S *in the direction of* $\mathbf{n}$. For example, if S is a segment of length 0.1 meters tilted at an angle of 45 degrees below the horizontal (so we can take $\mathbf{n} = (1/\sqrt{2}, 1/\sqrt{2})$) and $\mathbf{v} = (1, 2)$ meters per second then

$$|S|\mathbf{v}\cdot\mathbf{n} = 0.1(1/\sqrt{2} + 2/\sqrt{2}) \approx 0.212$$

square meters per second. The fact that the answer is positive indicates that the net flow is in the direction of the vector $\mathbf{n}$. Note that we could have taken the unit normal vector $\mathbf{n}$ as $(-\frac{1}{\sqrt{2}}, -\frac{1}{\sqrt{2}})$, in which case the answer would have been -0.212 square units per second, indicating the flow is opposed to $\mathbf{n}$.

One thing worth noting is this: the vector field $\mathbf{v}$, a velocity, can be thought of a bit differently. Obviously the direction of $\mathbf{v}$ is the direction of fluid flow. What about the magnitude of $\mathbf{v}$? Suppose that the segment S is orthogonal to the field $\mathbf{v}$, that is, $\mathbf{n}$ is parallel to $\mathbf{v}$. In this case the amount A of square units of fluid crossing S per unit time is given by $A = |S|\mathbf{v}\cdot\mathbf{n} = |S||\mathbf{v}|$ square units per second. This can be rearranged into $|\mathbf{v}| = A/|S|$, with $|\mathbf{v}|$ having the physical dimension of square units per second per unit length. Based on this, we can think of $\mathbf{v}$ as a vector which points in the direction of fluid flow, but with magnitude $|\mathbf{v}|$ given by the rate at which fluid is crossing a short segment S oriented orthogonal to the flow, per length of S. The advantage of thinking of $\mathbf{v}$ as "stuff per time crossing S per length of S" is that this interpretation extends to physical setting in which the "stuff" that flows doesn't really have a velocity, but where we can still think about how much stuff per time crosses a given curve (e.g., heat energy).

Now suppose that the curve is not a simple line segment, and the vector field $\mathbf{v}$ varies from point to point. How do we compute the amount of fluid crossing the curve per unit time? Let C be the curve, described parametrically by a function $\mathbf{r}(t) = (x(t), y(t))$ with $t_1 \leq t \leq t_2$. We can compute the fluid crossing C by chopping the curve C into many short segments. Over each segment the vector field $\mathbf{v}$ is nearly constant. We use the procedure above to find the fluid crossing each piece and then add them up.

The short pieces that we chop C into will be obtained by letting t change from t to $t + dt$; the resulting change in $\mathbf{r}(t)$ will be

$$d\mathbf{r} = \mathbf{r}(t + dt) - \mathbf{r}(t) \approx \mathbf{r}'(t)\, dt.$$

Here $d\mathbf{r}$ is vector that plays the role of S above. Its length is $|\mathbf{r}'|dt = ((x'(t))^2 + (y'(t))^2)^{1/2}dt$. We also need a unit normal vector to $d\mathbf{r} = (x'(t), y'(t))\, dt$. In two dimensions, a normal vector can be obtained by reversing the components of the vector and making one negative, e.g., the vector $(y'(t), -x'(t))$ is normal to $d\mathbf{r}$. To make it a unit vector, divide it by its own length, to obtain

$$\mathbf{n} = \frac{(y'(t), -x'(t))}{((x'(t))^2 + (y'(t))^2)^{1/2}}.$$

If the vector field is $\mathbf{v}(x,y) = (v_1(x,y), v_2(x,y))$ then the flux over this short piece is

$$d(flux) = \mathbf{v} \cdot \mathbf{n}|d\mathbf{r}| = (v_1(x,y)y'(t) - v_2(x,y)x'(t))\, dt$$

where the hideous square root in the denominator of $\mathbf{n}$ is cancelled by an identical square root in $d\mathbf{r}$. The total flux over the whole curve C is obtained by adding up over each piece,

$$\int_{t_1}^{t_2} (v_1(x,y)y'(t) - v_2(x,y)x'(t))\, dt. \tag{1}$$

Of course I could have chosen $\mathbf{n}$ to point the other direction. That's ok, you just have to keep track of what you choose and interpret the answer accordingly. The integral (1) is called the *flux* over the vector field $\mathbf{v}$ over the curve C.

Example: Let the curve C be the unit circle, $\mathbf{r}(t) = (\cos(t), \sin(t))$ for $0 \le t \le 2\pi$. It's worth noting that $\mathbf{n} = (\cos(t), \sin(t))$ (obvious—why?) and that $|d\mathbf{r}| = dt$, although given formula (1) we don't explicitly need these facts. Note that $\mathbf{n}$ points out of the circle. Let the vector field be $\mathbf{v}(x,y) = (x, yx^2)$. Then the amount of fluid crossing the boundary C per unit time is, from equation (1),

$$\int_0^{2\pi} (\mathbf{v} \cdot \mathbf{n})|d\mathbf{r}| = \int_0^{2\pi} (\sin^2(t)\cos^2(t) + \cos^2(t))\, dt = 5/4\pi$$

The answer is negative because the net flow of fluid over C is opposed to the outward-pointing vector $\mathbf{n}$, so that we have a net flow of $-5/4\pi$ square meters per second INTO the circle.

More generally, given any vector field $\mathbf{F}$ (not necessarily the velocity of any fluid) and a curve C we can talk abstractly about the flux of $\mathbf{F}$ over C. In two-dimensions the flux of $\mathbf{F}$ over a one dimensional curve C is written abstractly as

$$\int_C \mathbf{F} \cdot \mathbf{n}\, ds$$

where $ds = |d\mathbf{r}|$ and $\mathbf{n}$ is the outward unit normal vector on C. In order to actually compute the flux we must parameterize C and write out the integrand (this is formula (1), use $\mathbf{F}$ instead of $\mathbf{v}$). It is a fact that how one parameterizes the curve C does not affect the value of the flux integral!

0.7 Surface Integrals

This is a simple generalization of line integrals; only the details of the computation change. As an example, suppose we have a surface specified parametrically as

$$\begin{aligned} x &= \cos(v)(5 - 2\cos(u)), \\ y &= \sin(v)(5 - 2\cos(u)), \\ z &= 2\sin(u), \end{aligned}$$

for $0 \leq u, v \leq 2\pi$. (You can plot this in Maple—it's a torus.) Let's also use the notation $\mathbf{r}(u, v)$ when convenient, so $\mathbf{r}$ is a three-dimensional vector. Suppose that the density of the surface (mass per unit area) is given by

$$\rho(u, v) = 2 + u + v^2 - v$$

with units of kg per square meter. What's the mass of the surface? The idea is exactly the same as a line integral: chop the surface into many small pieces, compute the mass of each, add them up.

To chop up the surface we'll take small changes in u and v. Suppose that u changes to $u + du$. Then $\mathbf{r}$ changes by an amount $d\mathbf{r}_u$ with

$$d\mathbf{r}_u = \mathbf{r}(u + du) - \mathbf{r}(u) \approx (x_u, y_u, z_u)\, du$$

where I'm writing x_u for $\frac{\partial x}{\partial u}$, etc. The vector $d\mathbf{r}_u$ is an "infinitesimal" vector which is tangent to the surface $\mathbf{r}(u, v)$. Similarly, a small change in v to $v + dv$ sweeps out a vector

$$d\mathbf{r}_v = \mathbf{r}(v + dv) - \mathbf{r}(v) \approx (x_v, y_v, z_v)\, dv$$

which is also tangent to the surface. Consider the area spanned by these two vectors; it's a parallelogram. This will be our strategy for chopping up the surface. Perhaps it's best illustrated with a graphic like

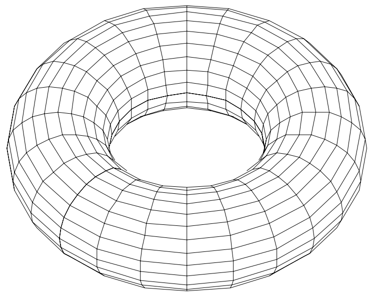

See how Maple draws the surface as a bunch of small quadrilaterals? We're doing the same thing, but using parallelograms (and in the limit that we chop finely there's really no difference). The corner of each parallelogram represents some parameter values (u, v), while the other 3 corners correspond to $(u + du, v)$, $(u, v + dv)$, and $(u + du, v + dv)$. The parallelogram's sides are spanned by $d\mathbf{r}_u$ and $d\mathbf{r}_v$. The area dA of the parallelogram is just (from high school algebra) $|d\mathbf{r}_u||d\mathbf{r}_v|\sin(\theta)$, where θ is the acute angle between the vectors. This is most easily found (in three dimensions) by using the cross product. You might recall that a basic property of the cross product is that

$$|d\mathbf{r}_u \times d\mathbf{r}_v| = |d\mathbf{r}_u||d\mathbf{r}_v|\sin(\theta)$$

so $dA = |d\mathbf{r}_u \times d\mathbf{r}_v|$ in this case. For our torus case the tedious computation of dA gives

$$dA = \sqrt{100 - 80\cos(u) + 16\cos^2(u)}\, du\, dv.$$

You can find the area of the torus by adding up all the dA's over the appropriate u and v limits,

$$\int_0^{2\pi}\int_0^{2\pi} \sqrt{100 - 80\cos(u) + 16\cos^2(u)}\, dv\, du = 40\pi^2.$$

The mass of any individual piece is $\rho\, dA$, so the total mass will be the sum $\int \rho\, dA$, which is

$$\int_0^{2\pi} \int_0^{2\pi} \sqrt{100 - 80\cos(u) + 16\cos^2(u)}(2 + u + v^2 - v)\, dv\, du = 80\pi^2 + \frac{160}{3}\pi^4$$

with units of kilograms.

0.8 Flux over a Surface in Three Dimensions

Conceptually, this is the same as the flux over a curve in two-dimensional space. First, suppose that L is a planar surface in three-dimensional space, with unit normal vector $\mathbf{n}$, and $\mathbf{F}$ is a constant vector field which represents the volume flow rate of some (incompressible) fluid. Note that $\mathbf{F}$ would have units of volume per unit time per AREA. How much fluid per unit time flows over the surface L? You should be able to convince yourself that it's just $|L||\mathbf{F}|\sin(\frac{\pi}{2} - \theta)$, where $|L|$ is the area of L, $|\mathbf{F}|$ is the length of $\mathbf{F}$ (the speed of the fluid), and θ the angle between the surface and $\mathbf{F}$. Of course, this is also $|L||\mathbf{F}|\cos(\theta)$ or just $|L|\mathbf{F} \cdot \mathbf{n}$. It's just like in two-dimensional space.

Now, what if $\mathbf{F}$ varies from point to point, and what if the surface isn't planar? You know the drill: Slice and dice. Let's denote the surface by S. We're going to chop S up into little pieces; each piece looks like a plane, has area dA, and has some unit normal vector $\mathbf{n}$. By the above reasoning, the flux over this little piece is then $\mathbf{F} \cdot \mathbf{n}\, dA$. The total flux of $\mathbf{F}$ over the surface is

$$\int_S \mathbf{F} \cdot \mathbf{n}\, dA. \tag{2}$$

Of course, we need to know how to actually compute this integral. In practice the surface would be specified parametrically, as $\mathbf{r}(u, v)$ where u and v range over some values. The field $\mathbf{F}$ would be given. We already saw how to cook up dA. It's just

$$dA = |\mathbf{r}_u \times \mathbf{r}_v|\, du\, dv.$$

All we need is a unit normal vector $\mathbf{n}$. Actually, we've already done this too. Look at how we got dA. We started with tangent vectors $\mathbf{r}_u$ and $\mathbf{r}_v$, took their cross product, then took the magnitude of the cross product. But the vector (note that I've written the A in boldface)

$$d\mathbf{A} = d\mathbf{r}_u \times d\mathbf{r}_v$$

is normal to the surface. We can define a unit normal vector by dividing $d\mathbf{A}$ by its own length, which is just dA:

$$\mathbf{n} = \frac{d\mathbf{A}}{|d\mathbf{A}|} = \frac{d\mathbf{A}}{dA}.$$

This shows how to actually compute the flux in a specific case. Compute dA and compute $\mathbf{n}$. Both should be functions of the independent parameters u and v. Stuff all this (with $\mathbf{F}$) into the flux integral (which will be a double integral in u and v) and evaluate.

There is one small simplification you can make. Note that $\mathbf{n}\, dA = d\mathbf{A}$, so that the flux integral can also be written as

$$\int_S \mathbf{F} \cdot d\mathbf{A} \tag{3}$$

which you'll frequently see.

Example: Let's use the torus defined previously. Let the vector field be $\mathbf{F}(x, y, z) = (xy, -z + 2, z + y)$. You can compute $d\mathbf{r}_u$ and $d\mathbf{r}_v$ to find

$$\begin{aligned} d\mathbf{r}_u &= (2\cos(v)\sin(u), 2\sin(v)\sin(u), 2\cos(u))\, du, \\ d\mathbf{r}_v &= (-\sin(v)(5 - 2\cos(u)), \cos(v)(5 - 2\cos(u)), 0)\, dv. \end{aligned}$$

Compute the vector $d\mathbf{A}$ as $d\mathbf{r}_u \times d\mathbf{r}_v$ (use a computer!),

$$\begin{aligned} d\mathbf{A} &= (-10\cos(u)\cos(v) + 4\cos^2(u)\cos(v), -10\cos(u)\sin(v) + 4\cos^2(u)\sin(v), \\ &\quad 10\sin(u) - 4\sin(u)\cos(u))\, du\, dv. \end{aligned}$$

Now compute $\mathbf{F} \cdot d\mathbf{A}$ (and replace the x,y, and z's in $\mathbf{F}$ with the corresponding u and v values). The quantity $\mathbf{F} \cdot d\mathbf{A}$ is a huge mess involving trig functions of u and v. Integrate it over the appropriate u and v ranges and you find that the total flux is $40\pi^2$.

0.9 Notation

It's frequently advantageous to think of the gradient operator as a vector, so that in two dimensions

$$\nabla = \left(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}\right).$$

Of course you can modify this to suit other dimensions. The operator ∇ sits around waiting for a function; when it encounters one, it computes its gradient.

Let $\mathbf{F}(x, y) = (F_1(x, y), F_2(x, y))$ be a vector field, and consider the quantity $\nabla \cdot \mathbf{F}$. What should that mean? If ∇ is treated as a vector then a logical interpretation is that

$$\nabla \cdot \mathbf{F} = \frac{\partial F_1}{\partial x} + \frac{\partial F_2}{\partial y}.$$

IMPORTANT: The above statement doesn't contain any real meaning or mathematics. It simply defines what we mean when we write $\nabla \cdot \mathbf{F}$. In other words, it's a definition. The quantity $\nabla \cdot \mathbf{F}$ is called the *divergence* of the vector field $\mathbf{F}$, for reasons that will become clear.

Now, suppose that the vector field $\mathbf{F}$ is itself the gradient of some function f, so $\mathbf{F} = \nabla f$. Then the divergence of $\mathbf{F}$ is

$$\nabla \cdot \nabla f = \frac{\partial^2 f}{\partial x^2} + \frac{\partial^2 f}{\partial y^2}.$$

The quantity $\nabla \cdot \nabla f$ is called the *Laplacian* of f. You'll also see it written as $\nabla^2 f$ (physicists like this) and as $\triangle f$, which is the mathematicians' favorite notation.

Given a vector field $\mathbf{F}$ you can also compute the quantity $\nabla \times \mathbf{F}$. I'll let you figure out the details. It's just the cross product. This quantity is called the *curl* of the vector field $\mathbf{F}$.

The ∇ notation can be a real time saver, once you get used to it. We are frequently going to manipulate expressions involving ∇, functions, and vector fields. For example, consider the expression $\nabla(\mathbf{F} \cdot \nabla f)$, where $\mathbf{F}$ is some vector field and f is some function. When you get used to the notation you can easily see that when expanded

$$\nabla(\mathbf{F} \cdot \nabla f) = (\nabla \cdot \mathbf{F})\nabla f + (\triangle f)\mathbf{F}.$$

The alternative is to write everything out component-by-component, which can be quite tedious. You should probably do this at first, but eventually you'll come to appreciate the brevity and power of the vector and ∇ notation, and you won't have to write every component out.

0.10 The Divergence Theorem

Suppose D is a region in 2 or 3 dimensions. We'll use the notation ∂D to denote the boundary of D. Suppose $\mathbf{F}$ is a vector field defined on D, with the first partial derivatives of $\mathbf{F}$ continuous on the closure of D. Then the divergence theorem states that

$$\int_D \nabla \cdot \mathbf{F}\, dV = \int_{\partial D} \mathbf{F} \cdot d\mathbf{A}, \tag{4}$$

where $d\mathbf{A}$ (or the unit normal vector $\mathbf{n}$) is chosen to point out of D.

Let's look at a concrete example or two, then a more general case to try to gain an understanding of what the divergence $\nabla \cdot \mathbf{F}$ tells us about a vector field.

Example 1: Suppose that $\mathbf{F}(x, y) = (x, y)$, and let D be the unit disk with boundary ∂D, the unit circle. Here's a picture of $\mathbf{F}$ superimposed over D.

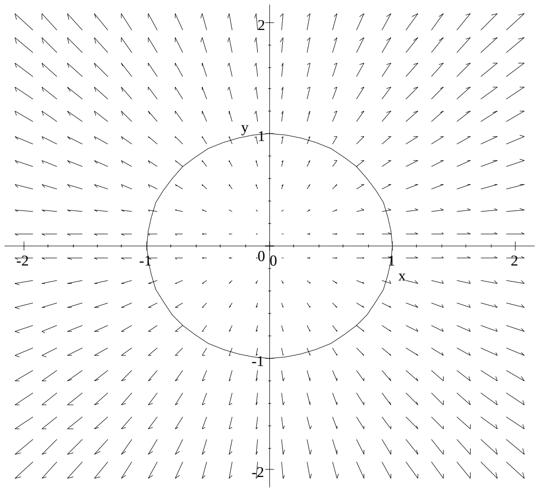

It should be painfully obvious that the outward flux of the field is positive. In fact, if we parameterize ∂D as $\mathbf{r}(t) = (\cos(t), \sin(t))$ for $0 \leq t < 2\pi$ then

$d\mathbf{r} = (-\sin(t), \cos(t))\,dt$, $ds = |d\mathbf{r}| = dt$, and an outward unit normal vector is given by $\mathbf{n} = (\cos(t), \sin(t))$. The flux over ∂D is

$$\int_{\partial D} \mathbf{F} \cdot \mathbf{n}\,ds = \int_0^{2\pi} (\cos^2(t) + \sin^2(t))\,dt = 2\pi.$$

It's easy to check that $\nabla \cdot \mathbf{F} = 2$. Then you can set up a double integral to compute

$$\int_D \nabla \cdot \mathbf{F}\,dy\,dx = \int_{-1}^{1} \int_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} 2\,dy\,dx = 2\pi$$

if you work it out. Of course, it must be the same as the flux of $\mathbf{F}$ over the boundary.

Example 2: Suppose that $\mathbf{F}(x, y) = (-y, x)$, and again let D be the unit disk with boundary ∂D, the unit circle. Here's a picture of $\mathbf{F}$ superimposed over D.

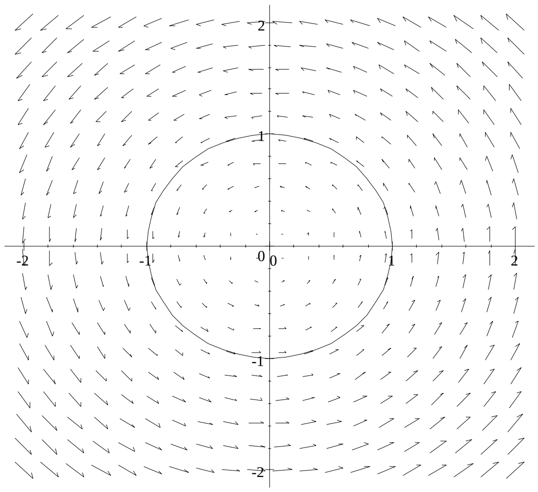

In this case the fluid spins tangentially to ∂D, so the net flux is zero. Let's do the computation anyway. We again parameterize ∂D as $\mathbf{r}(t) = (\cos(t), \sin(t))$ for $0 \le t < 2\pi$, with $d\mathbf{r} = (-\sin(t), \cos(t))\,dt$, $ds = |d\mathbf{r}| = dt$, and an outward unit normal vector is given by $\mathbf{n} = (\cos(t), \sin(t))$. The flux over ∂D

is

$$\int_{\partial D} \mathbf{F} \cdot \mathbf{n}\, ds = \int_0^{2\pi} (-\sin(t)\cos(t) + \sin(t)\cos(t))\, dt = 0$$

as asserted. It's easy to check that $\nabla \cdot \mathbf{F} = 0$. Obviously the integral of 0 over D is 0, consistent with the divergence theorem.

Let's now try to get a feeling for what $\nabla \cdot \mathbf{F}$ really means. In a nutshell, the divergence of a vector field measures how much the field "spreads out" near a point. To help you see this, suppose that $\mathbf{F}$ represents the flow (velocity) of some incompressible fluid in two or three dimensional space. Suppose P is a point in space and D is a small ball around P, sufficiently small that we can consider $\nabla \cdot \mathbf{F}$ constant on D. Then the flux of $\mathbf{F}$ over ∂D,

$$\int_{\partial D} \mathbf{F} \cdot d\mathbf{A}$$

is the rate at which fluid is crossing ∂D (with outward as positive) in area per second (2D) or volume per second (3D). According to the divergence theorem

$$\int_{\partial D} \mathbf{F} \cdot d\mathbf{A} = \int_D \nabla \cdot \mathbf{F}\, dV \approx (\nabla \cdot \mathbf{F}(P))\mathrm{vol}(D)$$

where $\mathrm{vol}(D)$ means the volume of D. In other words, for a small region D

$$\nabla \cdot \mathbf{F}(P) \approx \frac{\text{flux over } \partial D}{\mathrm{vol}(D)}$$

But why would there be a net flux over ∂D? Because inside D fluid is being created (if the outward flux is positive) or destroyed (outward flux negative). In three dimensions $\nabla \cdot \mathbf{F}$ has units of volume per unit time per volume. In two dimensions it has dimensions of area per second per unit area. The divergence of $\mathbf{F}$ at a point P in this case measures to what extent fluid is being created near P (divergence greater than zero) or fluid is being destroyed near P (divergence less than zero). For example, if in a region D the divergence of $\mathbf{F}$ is a constant 3 cubic meters per second per cubic meter then this means that during each second each cubic meter of volume in D produces 3 cubic meters of fluid.

Of course in real fluid flow the volume of fluid is conserved *if the fluid is incompressible*. In this case if $\mathbf{F}$ is the fluid velocity then $\mathbf{F}$ must have zero divergence, so that

$$\nabla \cdot \mathbf{F} = 0.$$

This is the so-called *continuity equation* for the case of steady state flow.

0.11 Green's Identities

Green's identities are really special cases of the divergence theorem that turn out to be very useful for studying PDE's. Suppose that u and v are C^2 functions defined on some region D in two or three (or more) dimensional space. Let ∂D denote the boundary of D; in two dimensions ∂D will be a closed curve, while in three dimensions ∂D is a closed surface. Consider the vector field $\mathbf{F} = u\nabla v$. According to the divergence theorem

$$\int_D \nabla \cdot (u\nabla v)\, dV = \int_{\partial D} (u\nabla v) \cdot \mathbf{n}\, dA \tag{5}$$

where $\mathbf{n}$ is an outward unit normal vector. But you can easily check (do it component by component—it's the product rule) that

$$\nabla \cdot (u\nabla v) = \nabla u \cdot \nabla v + u \triangle v$$

and

$$(u\nabla v) \cdot \mathbf{n} = u\frac{\partial v}{\partial \mathbf{n}}.$$

In this case equation (5) becomes

$$\int_D (\nabla u \cdot \nabla v + u \triangle v)\, dV = \int_{\partial D} u\frac{\partial v}{\partial \mathbf{n}}\, dA. \tag{6}$$

This is Green's first identity. A useful special case is when $u = v$, in which case we have

$$\int_D (|\nabla u|^2 + u \triangle u)\, dV = \int_{\partial D} u\frac{\partial u}{\partial \mathbf{n}}\, dA \tag{7}$$

where $|\nabla u|^2$ is just $\nabla u \cdot \nabla u$.

Take Green's first identity (6) and reverse the roles of u and v to obtain

$$\int_D (\nabla u \cdot \nabla v + v \triangle u)\, dV = \int_{\partial D} v\frac{\partial u}{\partial \mathbf{n}}\, dA. \tag{8}$$

Now subtract equation (8) from (6) to find that

$$\int_D (u \triangle v - v \triangle u)\, dV = \int_{\partial D} (u\frac{\partial v}{\partial \mathbf{n}} - v\frac{\partial u}{\partial \mathbf{n}})\, dA. \tag{9}$$

This is Green's second identity.

There's also a third identity, but we'll look at that later.

Example: Here's a simple example that's consistent with Green's first identity. Suppose that D is the unit ball in two dimensions and $u(x, y) = x^3 + 3x + xy$. Let's compute both sides of Green's identity and verify that they're the same. You can check that $\nabla u = (3x^2 + 3 + y, x)$ and $\triangle u = 6x$. The left side of Green's first identity, equation (6) (with appropriate limits on the double integral) is

$$\int_{-1}^{1} \int_{-\sqrt{1-x^2}}^{\sqrt{1-x^2}} (15x^4 + 37x^2 + 12x^2y + 9 + 6y + y^2)\, dy\, dx$$

which evaluates to $\frac{163}{8}\pi$. Now let's compute the right side. We can parameterize the boundary of the disk as $\mathbf{r}(t) = (\cos(t), \sin(t))$. Then $\mathbf{n} = (\cos(t), \sin(t))$ and $|d\mathbf{r}| = dt$. Then $u(\cos(t), \sin(t)) = \cos^3(t) + 3\cos(t) + \cos(t)\sin(t)$ on the boundary, and

$$\frac{\partial u}{\partial \mathbf{n}} = \nabla u \cdot \mathbf{n} = 3\cos^3(t) + 3\cos(t) + 2\sin(t)\cos(t)$$

and the product $u\frac{\partial u}{\partial \mathbf{n}}$ becomes, on the boundary,

$$u\frac{\partial u}{\partial \mathbf{n}} = 3\cos^6(t) + 12\cos^4(t) + 5\cos^4(t)\sin(t) + 9\cos^2(t)(1 + \sin(t)) + 2\cos^2(t)\sin^2(t).$$

If you integrate this from $t = 0$ to $t = 2\pi$ you get exactly $\frac{163}{8}\pi$!