

Quantitative Understanding in Biology

Module V: PDEs

Lecture II: Div, Grad, and Curl

Introduction

This lecture is a bit different from many others in this course in that it is intended as a survey of a topic, that of vector calculus. While most other lectures were aimed at giving you some practical skills to take away and use, the primary goal here is to develop (or reinforce) an appreciation for some of the basic principles in vector calculus, and how the topic is applied to biological problems. Another way of putting it, we seek to explain what all of those upside triangles are in some intimidating equations you may run across. You saw one such instance when Dr. Kressler introduced the Laplacian operator in the last lecture on image processing.

Vector calculus is, of course, a large and complex topic that can be the subject of a two semester course. Trying to cover everything in one and a half hours is clearly not feasible, and we won't try. In the interests of time, and understanding(!), much of what we will cover today will rely on intuition, and we won't be completely rigorous about every step. Much of this lecture is based on material in the excellent text (pamphlet) "Div, grad, curl and all that", by H. M. Schey. This 150 page easy-to-read book is one of my personal favorite math texts. It is easy to read, affordable (\$35), and should be in everyone's library.

Preliminaries

Before we dig into the details, we need to set up a few preliminary ideas and conventions. The first is that of a right-handed coordinate system. You probably recall something along these lines from physics. To check if a coordinate system is right-handed, begin by aligning your right thumb along the positive x-axis and your right index finger (pointer) along the positive y-axis. Finally, position your middle finger so that it is perpendicular to both; it will point along the positive z-axis. Most mathematicians use right-handed coordinate systems, but be warned that many computer graphics applications use a left-handed system. We'll use a right-handed system here. Incidentally, when using many graphics programs you can often add a set of axes to the screen to orient yourself; by convention the x-, y-, and z-axes are colored red, green, and blue, respectively.

Next, we should understand the concept of a scalar field and a vector field. A scalar field is a function of space where the function associates a scalar (i.e., a single number) with each point in space. An example of a scalar field is the temperature at every point in this room; another is the concentration of oxygen at every point. A vector field is a function where each point in space is associated with a vector; one example might be the velocity of air at each point in space in the room.

Div, Grad, and Curl

As you have probably guessed by now, we are going to be dealing with functions in space quite a bit today. We will be switching between functions in 2D space and 3D space pretty quickly and somewhat loosely as we progress. We often prefer to deal with the 2D cases because it is easy to draw pictures and visualize certain concepts. However, some topics in vector calculus only work in 3D space, so we will have to try to develop some agility in our spatial reasoning.

An alternative way of thinking about a vector field in \mathfrak{R}^3 (mathematicians' fancy way of referring to 3D space) is as a set of three scalar fields, each corresponding to one of the x-, y-, or z-components of the vector. We will use the notation F_x , F_y , and F_z to represent these components of the vector \mathbf{F} , so $F_x = \mathbf{F} \cdot \mathbf{i}$. We use \mathbf{i} , \mathbf{j} , and \mathbf{k} to denote the unit vectors along the three Cartesian axes throughout. Note that in some texts F_x is defined as $\frac{\partial F}{\partial x}$; we will *not* be using this convention.

At this point, we should define the cross product of two vectors in \mathfrak{R}^3 . The cross-product of \mathbf{u} and \mathbf{v} , written $\mathbf{u} \times \mathbf{v}$, is defined as...

$$\mathbf{u} \times \mathbf{v} = uv \sin \theta \hat{\mathbf{n}}$$


...where u is the norm of \mathbf{u} , v is the norm of \mathbf{v} , the angle between the vectors is θ , and $\hat{\mathbf{n}}$ is a unit vector perpendicular to both \mathbf{u} and \mathbf{v} , and chosen such that it points in the direction given by the right-hand rule. You can compute this numerically as

$$\mathbf{u} \times \mathbf{v} = (u_y v_z - u_z v_y)\mathbf{i} + (u_z v_x - u_x v_z)\mathbf{j} + (u_x v_y - u_y v_x)\mathbf{k}$$

...which is also written as...

$$\mathbf{u} \times \mathbf{v} = \det \begin{bmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ u_x & u_y & u_z \\ v_x & v_y & v_z \end{bmatrix}$$

You may recall the mnemonic for evaluating 3x3 determinants by augmenting the matrix and summing and subtracting diagonals...



$$\det \mathbf{A} = \sum (\text{solid diagonals}) - \sum (\text{dashed diagonals})$$

Our final preliminary comment has to do with the regions of space we are considering our scalar and vector fields to be defined in. We will only consider simply connected regions of space. A region is simply connected if any closed curve within the region can be shrunk to a point without leaving the region (you are allowed to move the curve as you shrink it). For example, the space between two concentric spheres

Div, Grad, and Curl

is simply connected, while the space occupied by a torus (donut) is not. Some of the conclusions we will draw will only apply to simply connected regions. We will also restrict ourselves to “smooth” scalar and vector fields; this implies all functions are continuous, differentiable, and have continuous first derivatives.

Surface Integrals and Flux

Consider a surface S . We often want to evaluate an integral of a scalar field over a surface. We write this as a double integral...

$$\iint_S F \, dS$$

The idea is that you take the surface S and divide it into a bunch of small patches. For each patch, you compute the value of F at the center of the patch and multiply that by the area of the patch. The value of the integral is the sum of these products as the patches become differentially small (and their number approaches infinity).

The area of any surface S is just $\iint_S dS$; this is the special case where $F(x,y,z) = 1$. If a surface were fabricated of a thin sheet of material with density $\rho(x,y,z)$ (in gm/cm^2), then the mass of the object would be $\iint_S \rho \, dS$.


The definition of a surface integral over a scalar field should be intuitive (luckily, we won't worry about how to compute them). We now consider what is usually implied when one considers the surface integral over a vector field. In fact, surface integrals are only defined on scalar fields, but when we consider vector fields we usually want the surface integral of the component of the vector fields that is perpendicular (aka 'normal') to the surface at any given point. In other words, we usually want...

$$\iint_S \mathbf{F} \cdot \hat{\mathbf{n}} \, dS$$

Remember that $\mathbf{F} \cdot \hat{\mathbf{n}}$ is the projection of \mathbf{F} onto $\hat{\mathbf{n}}$. So $\mathbf{F} \cdot \hat{\mathbf{n}}$ is a scalar that is defined by both the vector field (\mathbf{F}), and the surface (via $\hat{\mathbf{n}}$). Usually we have a mathematical function for \mathbf{F} and for S . So if we can figure out how to write $\hat{\mathbf{n}}$ in terms of S , we could in principle evaluate these surface integrals.

One way to get a vector normal to a surface is to generate two vectors tangent to the surface, and then take their cross product. Since the cross product is perpendicular to both vectors, it will be normal to the surface at that point. We'll assume here that our surface can be expressed as $z = f(x,y)$. While any two (non-colinear) vectors tangent to S will do, it is easiest to consider one in the x -direction, and one in the y -direction. For the x -direction case, consider the plane parallel to the x - z plane that cuts through our surface at the point of interest. The intersection of this plane and the surface is a curve, as shown here...

Div, Grad, and Curl



A vector tangent to this curve is also tangent to the surface S . We can generate such a vector by choosing an arbitrary distance along the x -axis, u_x , for the x -component of the tangent vector. Now the y -component of the tangent vector is given by $u_x \frac{\partial z}{\partial x}$. Using these components, we can construct our first vector, \mathbf{u} , that is tangent to the surface...

$$\mathbf{u} = u_x \mathbf{i} + \frac{\partial z}{\partial x} u_x \mathbf{k} = u_x \left[\mathbf{i} + \frac{\partial z}{\partial x} \mathbf{k} \right]$$

Similar reasoning for a second cutting plane parallel to the y - z plane yields...

$$\mathbf{v} = v_y \mathbf{j} + \frac{\partial z}{\partial y} v_y \mathbf{k} = v_y \left[\mathbf{j} + \frac{\partial z}{\partial y} \mathbf{k} \right]$$

Now the cross-product of these two vectors (we drop the u_x and v_y since all we care about is the direction)...

$$\mathbf{u} \times \mathbf{v} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ 1 & 0 & \frac{\partial z}{\partial x} \\ 0 & 1 & \frac{\partial z}{\partial y} \end{vmatrix} = -\mathbf{i} \frac{\partial z}{\partial x} - \mathbf{j} \frac{\partial z}{\partial y} + \mathbf{k}$$

Finally, to get the *unit* normal vector, we must normalize this...

$$\hat{\mathbf{n}} = \frac{-\mathbf{i} \frac{\partial z}{\partial x} - \mathbf{j} \frac{\partial z}{\partial y} + \mathbf{k}}{\sqrt{\left(\frac{\partial z}{\partial x}\right)^2 + \left(\frac{\partial z}{\partial y}\right)^2 + 1}}$$


This looks a bit intimidating, but you should now understand where all of the pieces come from. Furthermore, if you also have a function for \mathbf{F} , you can write down and expand $\mathbf{F} \cdot \hat{\mathbf{n}}$, and, in principle, compute the integral.

Div, Grad, and Curl

These kinds of integrals are really important, If $\mathbf{F}(x,y,z)$ is a flow field (i.e., giving velocity) of some 'stuff' with density $\rho(x,y,z)$, then the integral...

$$\iint_S \rho \mathbf{F} \cdot \hat{\mathbf{n}} dS$$

...tells you how much stuff is flowing through the surface. This is called the flux (of 'stuff') through S . As an example, consider a uniform flow field of uniform density ρ through a rectangular channel (the fluid would have to be inviscid for this to be realistic, but don't worry too much about that) as shown in the figure below. The channel's height (into the page) will be taken to be H (not shown), and its width is W .



In the figure above, it is pretty easy to compute the flux through S_1 . The vector normal to the surface, $\hat{\mathbf{n}}_1$, is simply \mathbf{i} . \mathbf{F} is simply $v_x \mathbf{i}$. So $\mathbf{F} \cdot \hat{\mathbf{n}}_1 = v_x$. This is a constant and, along with ρ , can come out of the integral, leaving the integral to be simply the area of the surface, which is WH . So the flux through the surface is $v_x \rho WH$. Note that the dimension of this result is mass/time.

The flux through S_2 is only a little bit more involved to compute. First, we note that the surface area of S_2 is a bit bigger than S_1 . However, the unit normal is a bit off of the horizontal, so when you take the projection of the flow vector \mathbf{F} onto $\hat{\mathbf{n}}_2$, the resultant scalar will be a bit smaller than v_x . These two effects will exactly cancel, and you'll get the exact same flux as before. **Exercise: assume S_2 is inclined at an angle θ , and prove this.** This, of course, has to be the case if our interpretation of flux is correct. The same holds for an arbitrarily shaped surface such as S_3 ; you'd need some real vector calculus skills to actually do the integral.

Consider the surfaces S_4 and S_5 . The flux through each of these surfaces is zero. Intuitively, we see that a lot of material is flowing around (or by) the surfaces, but nothing is flowing through them. Mathematically, the projection of the flow vector onto the surface normal is zero (because they are orthogonal), so the flux must be zero.

The flux through S_6 gives us a bit of pause, because we didn't consider one important aspect of the unit normal vector. For any given surface, there are two possible normals; and they point in opposite directions. For our S_1 case, we implicitly assumed that $\hat{\mathbf{n}}_1$ pointed to the right, but we could just as easily defined it as pointing to the left, in which case the flux would have been computed as $-v_x \rho WH$. For an arbitrary surface the choice of the two possible unit normals is arbitrary, and should be stated as part of the solution. However, in the case of closed surfaces, such as S_6 , which we will take to be capped on the

Div, Grad, and Curl

sides so it forms a cube, there is an accepted convention that the unit normal points out of the closed surface. Give this convention, you can reason out that the flux through S_6 is zero.

The Divergence

Let us consider the flux through a cube's surface such as S_6 a bit more closely, but this time in an arbitrary vector field instead of a simple uniform one. One natural way of characterizing the vector field at a point in space might be to consider what the flux is through a cube's surfaces as it gets differentially smaller. Sadly, the answer is unhelpful; it is zero. The surface area of each face of the cube goes to zero as the cube gets differentially smaller, and the components of the vector field are finite, so the product tends to zero. However, the flux per unit volume through the cube's faces will not tend to zero, and will yield a useful quantity.

Let us consider just the left and right faces of a small cube with dimensions Δx , Δy , and Δz . The flux through the right face is...

$$Flux_R = F_x \left(x + \frac{\Delta x}{2}, y, z \right) \Delta y \Delta z$$

...and the flux through the left face is

$$-Flux_L = F_x \left(x - \frac{\Delta x}{2}, y, z \right) \Delta y \Delta z$$

The net flux through the two faces is then...

$$Flux_{L+R} = F_x \left(x + \frac{\Delta x}{2}, y, z \right) \Delta y \Delta z - F_x \left(x - \frac{\Delta x}{2}, y, z \right) \Delta y \Delta z$$

Since the volume of the cube is $\Delta x \Delta y \Delta z$, the flux per unit volume is...

$$\frac{Flux_{L+R}}{\Delta V} = \frac{F_x \left(x + \frac{\Delta x}{2}, y, z \right) - F_x \left(x - \frac{\Delta x}{2}, y, z \right)}{\Delta x}$$

In the limit as the volume, and thus Δx , approaches zero, the right hand side of the above expression becomes the partial derivative of F_x with respect to x . So we have...

$$\lim_{\Delta V \rightarrow 0} \frac{Flux_{L+R}}{\Delta V} = \frac{\partial F_x}{\partial x}$$

The above argument only considered two of the six sides of the cube. If we considered all six sides, we would get the relation...

$$\lim_{\Delta V \rightarrow 0} \frac{Flux_{Total}}{\Delta V} = \frac{\partial F_x}{\partial x} + \frac{\partial F_y}{\partial y} + \frac{\partial F_z}{\partial z}$$

Div, Grad, and Curl

This quantity is a very useful characterization of the vector field at a single point. It is called the divergence of the vector field, as is sometimes written as $\text{div } \mathbf{F}$. However, it is more impressive to write...

$$\text{div } \mathbf{F} \stackrel{\text{def}}{=} \lim_{\Delta V \rightarrow 0} \frac{\text{Flux}_{\text{Total}}}{\Delta V} = \frac{\partial F_x}{\partial x} + \frac{\partial F_y}{\partial y} + \frac{\partial F_z}{\partial z} = \nabla \cdot \mathbf{F}$$

Implicit in this form of expressing the divergence is the notion of ∇ , which is pronounced “del”, as a mathematical object called an operator. Specifically, we can write...

$$\nabla = \mathbf{i} \frac{\partial}{\partial x} + \mathbf{j} \frac{\partial}{\partial y} + \mathbf{k} \frac{\partial}{\partial z}$$

If we recall that a vector function \mathbf{F} can be written as...

$$\mathbf{F} = \mathbf{i}F_x + \mathbf{j}F_y + \mathbf{k}F_z$$

...we can see how $\nabla \cdot \mathbf{F}$ leads to the result above.

If you think of your vector field as a flow, the divergence at a point in the field tells you if the point is acting as a source or a sink for the stuff that is flowing. A positive divergence indicates a net outflow of stuff from the point, and a negative divergence indicates a net inflow. A zero divergence indicates that stuff is just passing through (although it may be changing direction).

In the context of an electrical field, the divergence at a point tells you how much net charge there is at a point. If it is zero, field lines enter and leave the point in the same amount. If there is charge density, the divergence will be non-zero. Even though nothing is actually flowing in an electrostatic field, the analogy is often useful when visualizing fields.

Line Integrals and Work

We’ve covered integrals over surfaces in the previous section, and here we’ll cover something perhaps a little bit simpler; the line integral. The line integral is written as...

$$\int_C F ds$$

This is an integration along a parameterized curve, C . Imagine breaking the curve into a bunch of line segments. For each segment, compute the value of the scalar function F at the midpoint of the segment, and multiply it by the length of the segment. The value of the integral is the sum of all of the products so computed as the length of the segments tends to zero and the number of them tends to infinity. When we say that C is a parameterized curve, we mean that the x , y , and z points along the curve can be expressed as a function of the arc length along the curve...

$$x=x(s), y=y(s), z=z(s)$$

Div, Grad, and Curl

If we have a vector field, we do something very similar to the case for surface integrals; we compute the dot product of the vector function with an appropriate unit vector, and integrate the resultant scalar over the curve. In this case, the appropriate unit vector is the one tangent to (and in the direction of) the curve C . Such an integral is written as

$$\int_C \mathbf{F} \cdot \hat{\mathbf{t}} ds$$

In the interests of time (and space), we won't derive the general expression for $\hat{\mathbf{t}}$, and simply state...

$$\hat{\mathbf{t}} = \mathbf{i} \frac{dx}{ds} + \mathbf{j} \frac{dy}{ds} + \mathbf{k} \frac{dz}{ds}$$

One of the most fundamental line integrals is that for work. You may recall from physics that work is the product of force times displacement, often written as $W = F\Delta x$. A more general way of writing this is...

$$Work = \int_C \mathbf{F} \cdot \hat{\mathbf{t}} ds$$

This says that as an object moves through space along a curve C , the work done by the force field, \mathbf{F} , is the sum (integral) of the projection of the force field onto the direction of motion.

Consider the work done by a gravitational field on a (circularly) orbiting satellite. Since the orbit is circular, the motion is always perpendicular to the gravitational field. Therefore, the projection is zero, and we conclude no work is done. On the other hand, if we are launching a rocket, the motion is opposite to the force, so the work done is negative (i.e., we have to fight the field), and quite large.

Some force fields have a nifty property: the work done along the path from point A to point B is a function only of the beginning and ending points. In this case, we say that the line integral is path independent.

$$\int_A^B \mathbf{F} \cdot \hat{\mathbf{t}} ds = f(A, B)$$

A key implication of such a condition is that the integral around any closed path is zero, which we write as...


$$\oint_C \mathbf{F} \cdot \hat{\mathbf{t}} ds = 0$$

Forces that act along a vector connecting two particles, and whose magnitude is a function of the distance between those two particles, are called central forces, and have this mathematical property.

Div, Grad, and Curl

The Curl

In a manner analogous to the way we treated surface integrals to lead us to the divergence, we will consider a line integral around the girth of a differential volume centered on (x, y, z) . There are three such rings, and we'll begin by considering the one parallel to the top and bottom faces.



As with our approach to the divergence, the value of this line integral will tend towards zero as the volume becomes differentially small, but the ratio of the value of the line integral to the volume of the cube will converge on a fixed value.

$$\begin{aligned} \oint_C \mathbf{F} \cdot \hat{\mathbf{t}} \, ds &= \int_{C_B} F_x \left(x, y - \frac{\Delta y}{2}, z \right) ds_B - \int_{C_T} F_x \left(x, y + \frac{\Delta y}{2}, z \right) ds_T + \int_{C_R} F_y \left(x + \frac{\Delta x}{2}, y, z \right) ds_R \\ &\quad - \int_{C_L} F_y \left(x - \frac{\Delta x}{2}, y, z \right) ds_L \\ \oint_C \mathbf{F} \cdot \hat{\mathbf{t}} \, ds &= F_x \left(x, y - \frac{\Delta y}{2}, z \right) \Delta x - F_x \left(x, y + \frac{\Delta y}{2}, z \right) \Delta x + F_y \left(x + \frac{\Delta x}{2}, y, z \right) \Delta y \\ &\quad - F_y \left(x - \frac{\Delta x}{2}, y, z \right) \Delta y \\ \frac{1}{\Delta S} \oint_C \mathbf{F} \cdot \hat{\mathbf{t}} \, ds &= \frac{F_y \left(x + \frac{\Delta x}{2}, y, z \right) - F_y \left(x - \frac{\Delta x}{2}, y, z \right)}{\Delta x} - \frac{F_x \left(x, y + \frac{\Delta y}{2}, z \right) - F_x \left(x, y - \frac{\Delta y}{2}, z \right)}{\Delta y} \end{aligned}$$

To get the last line we have used the fact that $\Delta S = \Delta x \Delta y$. Now as we take the limit as the volume becomes differentially small, we have...

$$\lim_{\Delta S \rightarrow 0} \frac{1}{\Delta S} \oint_C \mathbf{F} \cdot \hat{\mathbf{t}} \, ds = \frac{\partial F_y}{\partial x} - \frac{\partial F_x}{\partial y}$$

This covers what we need for the loop normal to \mathbf{k} . We could do the same thing for each of the other two loops. This will give us three quantities, each of which can be associated with a vector normal to the surface formed by the loop. The final result is the curl, which we write as...

$$\nabla \times \mathbf{F} = \mathbf{i} \left(\frac{\partial F_z}{\partial y} - \frac{\partial F_y}{\partial z} \right) + \mathbf{j} \left(\frac{\partial F_x}{\partial z} - \frac{\partial F_z}{\partial x} \right) + \mathbf{k} \left(\frac{\partial F_y}{\partial x} - \frac{\partial F_x}{\partial y} \right)$$


Div, Grad, and Curl

This is a bit more easily remembered using the determinant form of the cross product...

$$\nabla \times \mathbf{F} = \begin{vmatrix} i & j & k \\ \frac{\partial}{\partial x} & \frac{\partial}{\partial y} & \frac{\partial}{\partial z} \\ F_x & F_y & F_z \end{vmatrix}$$

So the curl of a vector field is another vector field. The curl at a point characterizes a vector field by telling you how much 'swirliness' or net circulation there is at a given point. If you consider your vector field to be a flow, you can imagine a paddle-wheel embedded in the vector field. When the shaft of the paddle wheel points in the direction of the curl, the paddle wheel rotates most freely. Alternatively, to get a component of the curl, image a paddle wheel embedded in the flow with its axis pointing along the appropriate Cartesian axis. The speed at which the wheel would turn (at least initially, before the paddles changed the flow field) would give you the magnitude of the curl along that axis.


In the uniform channel flow field we considered above, the curl is zero. In the field below you might think there is no circulation, but in fact the curl is non-zero...



To see why, consider adding a paddle wheel into the field as shown. The force on vanes B and D are zero, and the force on A is in the positive x-direction while the force on C is in the negative x-direction. However, the force on A is larger than the force on C, so the wheel will have a tendency to start rotating about the shaft; therefore, this field has a curl.

Earlier we mentioned that central forces give rise to fields whose work function is path independent; here we will show that path independence of the line integral implies zero curl. While our level of rigor will somewhat less than what is needed for a formal proof, we will make what is hopefully a strong argument for this assertion.

Div, Grad, and Curl



Consider the two paths from 1 to 2 shown in the figure above. If the line integral is path independent, then we have...

$$\int_{L+T} \mathbf{F} \cdot \hat{\mathbf{t}} ds - \int_{B+R} \mathbf{F} \cdot \hat{\mathbf{t}} ds = 0$$

$$\int_{L+T} \mathbf{F} \cdot \hat{\mathbf{t}} ds = \int_{B+R} \mathbf{F} \cdot \hat{\mathbf{t}} ds$$

Using some of our relations from above, we have...

$$\int_{C_L} F_y \left(x - \frac{\Delta x}{2}, y, z \right) ds_L + \int_{C_T} F_x \left(x, y + \frac{\Delta y}{2}, z \right) ds_T =$$

$$\int_{C_B} F_x \left(x, y - \frac{\Delta y}{2}, z \right) ds_B + \int_{C_R} F_y \left(x + \frac{\Delta x}{2}, y, z \right) ds_R$$

...which yields...

$$F_y \left(x - \frac{\Delta x}{2}, y, z \right) \Delta y + F_x \left(x, y + \frac{\Delta y}{2}, z \right) \Delta x =$$

$$F_x \left(x, y - \frac{\Delta y}{2}, z \right) \Delta x + F_y \left(x + \frac{\Delta x}{2}, y, z \right) \Delta y$$

...or...

$$F_x \left(x, y + \frac{\Delta y}{2}, z \right) \Delta x - F_x \left(x, y - \frac{\Delta y}{2}, z \right) \Delta x =$$

$$F_y \left(x + \frac{\Delta x}{2}, y, z \right) \Delta y - F_y \left(x - \frac{\Delta x}{2}, y, z \right) \Delta y$$


Again considering a differentially small volume yields...

$$\frac{\partial F_x}{\partial y} = \frac{\partial F_y}{\partial x} \Rightarrow \frac{\partial F_y}{\partial x} - \frac{\partial F_x}{\partial y} = 0 \Rightarrow (\nabla \times \mathbf{F}) \cdot \mathbf{k} = 0$$

Div, Grad, and Curl

A similar argument can be made for the other two loops to consider. Overall, we get that path independence implies zero curl. It turns out that the reverse is also true, zero curl implies path independence (all of this with the provisos given in the Preliminaries section; in particular this only applies for simply connected spaces).

This is a very profound result, for it means that we can look at a field, and, by evaluating the curl everywhere, tell if the field could possibly be due to central forces (i.e., if it could be a gravitational field or an electrostatic field). For example, this field cannot be a gravitational field.



If it were, you would get some Escheresque work functions (reminiscent of the lines you may have heard from your grandparents – “when we were your age, we had to walk barefoot to school, *uphill, both ways!*”)

The Gradient

It turns out that the curl is most useful to us when it isn't there. Any point in a curl-free vector field can be characterized by a scalar whose value is the value of any line integral from an arbitrarily chosen reference point. We will call the resultant scalar field ψ . We have

$$\psi(x, y, z) = \int_{(x_0, y_0, z_0)}^{(x, y, z)} \mathbf{F} \cdot \hat{\mathbf{t}} \, dS + \psi_0$$

Note that this only makes sense if the line integral is path independent, otherwise, the value of ψ at a point would depend on the path to it from the reference point.

It should be clear that...

$$\frac{\partial \psi}{\partial x} = F_x \quad \frac{\partial \psi}{\partial y} = F_y \quad \frac{\partial \psi}{\partial z} = F_z$$

...or, written more compactly...

$$\mathbf{F} = \nabla \psi$$

Div, Grad, and Curl

To see this, consider the change in the value of ψ as we move from (x,y,z) to $(x+\Delta,y,z)$. We then have...

$$\psi(x + \Delta x, y, z) = \int_{(x,y,z)}^{(x+\Delta x,y,z)} \mathbf{F} \cdot \hat{\mathbf{t}} ds + \psi(x, y, z)$$

The value of the integral is $F_x \Delta x$, for small Δx , so we have...

$$\psi(x + \Delta x, y, z) - \psi(x, y, z) = F_x \Delta x$$

Dividing by Δx and taking the limit as it becomes differentially smaller gives us the first result above.

We say that ψ is the gradient of \mathbf{F} . We really like this result because it means that when we are working with central forces, we can think in terms of the scalar function $\psi(x, y, z)$ and not a more complex vector function $\mathbf{F}(x,y,z)$. ψ is the potential function, and its gradient is the force field.

Here we have introduced the gradient as a mathematical convenience for expressing a fields in terms of potentials. There is also an important geometric interpretation of the gradient, which we will now investigate.

Consider a function of two variables $z=f(x,y)$. We take the value of the function at the point (x_0, y_0) to be z_0 . Now consider that we move away from (x_0, y_0) in some arbitrary direction $\hat{\mathbf{u}}$. We wish to know how quickly z changes as we move in this direction. In other words, we want to know $\frac{\partial f}{\partial s}$ where s is in the direction of $\hat{\mathbf{u}}$. To tackle this problem, we will consider a Taylor series around (x_0, y_0) .

$$z \cong f(x_0, y_0) + \frac{\partial f}{\partial x} \Delta x + \frac{\partial f}{\partial y} \Delta y$$

$$\Delta f \cong \frac{\partial f}{\partial x} \Delta x + \frac{\partial f}{\partial y} \Delta y$$

We also have that...

$$\Delta \mathbf{s} = \mathbf{i} \Delta x + \mathbf{j} \Delta y$$

...and the definition of the gradient...

$$\nabla f = \mathbf{i} \frac{\partial f}{\partial x} + \mathbf{j} \frac{\partial f}{\partial y}$$

With these two relations in hand, we can rewrite the right hand side of the expression for Δf ...

$$\Delta f \cong \Delta \mathbf{s} \cdot \nabla f$$

We recall that $\Delta \mathbf{s}$ is in the direction of $\hat{\mathbf{u}}$, so we can write...

$$\Delta f \cong \Delta s (\hat{\mathbf{u}} \cdot \nabla f)$$

Div, Grad, and Curl

$$\frac{\Delta f}{\Delta s} \cong \hat{\mathbf{u}} \cdot \nabla f$$

...and then take the limit as the step becomes differentially small...

$$\frac{\partial f}{\partial s} = \hat{\mathbf{u}} \cdot \nabla f$$

In words, this says that the rate at which our function increases (or our surface rises) is a projection of the gradient onto the unit vector in the direction that we are moving.

If we want to maximize the rate of increase, we would choose a direction $\hat{\mathbf{u}}$ such that it is in the same direction as the gradient (recall that $\mathbf{a} \cdot \mathbf{b} = |\mathbf{a}||\mathbf{b}| \cos \theta$). This gives us a nice geometrical interpretation of the gradient of a scalar field; i.e., the gradient points in the direction of the steepest increase in f . You can now why we often make an analogy of potential fields as surfaces (height fields) and forces being akin to balls rolling down the hills of the potential surface (since the force is in the direction of the gradient).

The Laplacian

We will now revisit the divergence, but instead of writing it in terms of the field, \mathbf{F} , we will write it in terms of the potential, ψ .

$$\nabla \cdot \mathbf{F} = \nabla \cdot (\nabla \psi) = \left(\mathbf{i} \frac{\partial}{\partial x} + \mathbf{j} \frac{\partial}{\partial y} + \mathbf{k} \frac{\partial}{\partial z} \right) \cdot \left(\mathbf{i} \frac{\partial \psi}{\partial x} + \mathbf{j} \frac{\partial \psi}{\partial y} + \mathbf{k} \frac{\partial \psi}{\partial z} \right) = \frac{\partial^2 \psi}{\partial x^2} + \frac{\partial^2 \psi}{\partial y^2} + \frac{\partial^2 \psi}{\partial z^2} = \nabla^2 \psi$$

This sum of second derivatives of the potential field is called its Laplacian. It is nothing more than the divergence of the vector field corresponding to the potential.

A bit further back we alluded to Gauss' law, which states that the divergence of an electric field at each point is proportional to the charge density at that point. This can now be written in terms of the potential:

$$\nabla \cdot \mathbf{F} = \nabla^2 \psi = -\frac{\rho}{D\epsilon_0}$$

This is known as Poisson's equation, and is one of the fundamental equations of electrostatics (here we introduced D , which is the dielectric constant). A special case of this is the Laplace equation, which is just the special case that results when $\rho=0$. This is of interest when you have a fixed charge distribution and you want to compute the potential due to it in the charge-free spaces between the known charge distributions.

The Poisson-Boltzman Equation

An important equation in molecular biophysics is the Poisson-Boltzman equation. We begin by considering a fixed charge distribution, such as that due to a protein or DNA. In the case of DNA, the

Div, Grad, and Curl

molecule itself is charged (with negative phosphate groups). Usually, this will be offset by a stoichiometric amount of counterions. The counterions will tend to build up near the surface of the DNA, and effectively shield it from the rest of the solvent. If all you have is negatively charged DNA and mobile positive counterions, it becomes very difficult for two DNAs to come together because their shields of counterions will repel each other.

If, however, you add sufficient salt (which will result in both free positive and free negative ions), there is a lot more freedom in the way the system can arrange itself. Small, mobile negative ions can intervene in the shield layer, allowing the DNAs to come closer together. This opens the possibility for other short range interactions to take place, and can cause the DNA to precipitate.

One example of this phenomenon is river deltas. Small particles of silt tend to be charged, and, in fresh water, are not likely to precipitate. When a fresh water river reaches the ocean and the salinity of the water increases, the salt effect allows the particles to coalesce and precipitate, thus forming a river delta (from electrostatics to ecology – now that's multi-scale modeling for you!)

It is effects like this that make understanding the distribution of ions around a charged molecule important. Note that counterions are modeled as a continuum, not as discrete entities; so we are interested in computing the charge density, ρ , due to the ionic distribution. The challenge is that the ions don't just feel the electrostatic potential; they are part of the cause of it as well.

We already have half of the story; the Poisson equation above. The second half of the story is the Boltzman distribution. For positive ions, we can write what is essentially the Nernst equation...

$$n_+(x, y, z) = n_\infty e^{-\frac{ze\psi(x,y,z)}{kT}}$$

This says that positive ions like areas of negative potential, and abhor areas of positive potential. You can see this because when ψ is large, the exponential is small, and vice versa. The degree of separation is inversely proportional to temperature which makes sense too. When T is large, there are lots of random thermal fluctuations in the solutions, and the distribution is more even because of this. At low temperatures, everything finds its place of minimum energy, and at high temperatures, everything is all mixed up.

We can write a similar equation for negative ions...

$$n_-(x, y, z) = n_\infty e^{+\frac{ze\psi(x,y,z)}{kT}}$$

Charge density is then given by...

$$\rho(x, y, z) = ze[n_-(x, y, z) - n_+(x, y, z)]$$

Note that here we assume only two ionic species. In general, you may have more terms for additional ions (e.g., calcium, sodium, chloride, etc.)

Substituting this into the Poisson equation yields...

Div, Grad, and Curl

$$\nabla^2\psi = -\frac{zen_{\infty}}{D\epsilon_0}\left(e^{\frac{ze\psi}{kT}} - e^{-\frac{ze\psi}{kT}}\right)$$

This is a second order differential equation in ψ . The boundary conditions near the protein surface are complex, but the equation can be solved numerically on a computer.

One case that can be solved analytically is for that of an infinite charged plane in a salt solution. With no salt, the potential decays linearly as you move away from the plane. With salt, the decay is exponential. The 'length constant' of this decay, known as the Debye length, is proportional to the square root of n_{∞} .

More complex cases need to be solved numerically. We'll close by observing that a fair amount of supercomputer time is spent solving the equation above.