Volume Integrals

January 23, 2017

Overview

These notes cover the following:

- Densities as quantities defined 'per unit volume'
- 'Local' proportionality revisited: density as a 'field'
- Volume integrals
- Visualizing scalar fields

Density

Even though mass, and therefore other physically relevant quantites like momentum, energy, charge etc, are concentrated into particles like molecules that one could model as being separate at an atomic level or even subatomic level, this is frequently not a useful approach. In practical applications in geophysics, atmospheric science and many other areas, this would involve having to keep track of an extremely large number of particles, and actual computations would become impossible.

The study of *continua* is a way around this. In a continuum,[1] we assume that mass is spread out in space, rather than concentrated into particles. You can take this to be the definiton of a continuum in physics. A single point in a continuum therefore has no mass associated with it, but a finite volume can. To capture this mathematically, we need a way of describing how mass is distributed in space. The way to do this is using the idea of density, but reformulated as a *field*. By 'field', we simply mean a physically-relevant quantity that is a function of position and possibly of time.

[1] 'Continua' is plural of 'continuum'

You will know density as 'mass over volume',

$$\rho = \frac{m}{V}.$$

How can we make sense of this when we treat it as a field, so the density $\rho(x, y, z, t)$ is a function of position? A point (x, y, z) has no mass or volume, so we cannot take the mass of a point and divide it by the volume of the point.

The reason why we define density as mass over volume is that, for a material with spatially uniform properties, we expect mass to be proportional to volume: take twice the volume of iron, and you will get twice the mass, provided you keep temperature, pressure etc. constant. In other words, we have

$$m \propto V$$

and density ρ is the constant of proportionality.

This works well only for a material with uniform properties, and we are specifically interested in the case where there may be variations in space as well as time. The point is that we may expect these properties to change slowly enough in space that they can be treated as constant if we look at a very small volume. In that case, the mass δm in such a small volume δV should be proportional to the size of that volume

$$\delta m \propto \delta V$$

The constant of proportionality then depends on where that small volume is located, for instance on the location of its centre (x, y, z): it is a 'constant' with respect to the volume size δV (so long as δV is small) but not with respect to where that volume is located. Defining the constant of proportionality again as 'density', density becomes a function $\rho(x, y, z, t)$. We can define ρ through

$$\rho(x, y, z, t) = \frac{\text{mass } \delta m \text{ in a small volume } \delta V \text{ centered on } (x, y, z) \text{ at time } t}{\delta V},$$

or $\rho = \delta m / \delta V$ for short.

Note 1 *By 'small volume', we have to mean not only that the size of the volume is small, but that it is small in all directions. We have previously had an example of a volume that had small size but was not small in all directions: if you take a thin slice but keep the base area of the slice constant, then only the thickness of the slice shrinks as the size of the volume is made small, while the lateral extent remains the same. This is not what we mean by 'small volume'. The simples example of a small volume of the kind we have in mind would be a cuboid with three short side lengths δx, δy, δz.*

Volume integrals

If we know the density field $\rho(x, y, z, t)$, how do we go backwards: how do we calculate the mass M of a body occupying a known volume V? If V is large, we cannot simply compute mass as volume times density, since density varies over the volume. However, if we split the total volume into small pieces or 'volume elements' of size δV, then we can treat the density in each volume element as being approximately constant. This means we calculate the mass of each piece as

$$\delta m \approx \rho \delta V,$$

and sum over the masses of all the small pieces that make up the whole volume V:

$$M \approx \sum_{\delta V} \delta m \approx \sum_{\delta V} \rho \delta V.$$

Naturally, we expect the approximation invovled to become better and better as we make the division of V into small elements δV finer and finer. In other words, we have something like a Riemann sum, but over small volumes δV rather than over short intervals δx or δt. In the limit, we denote the Riemann sum as a 'volume integral', written in the form

$$M = \int_V \rho \, \mathrm{d}V,$$

where the subscript 'V' indicates the volume the integral is taken over.

This is all very well, but we do not have a way of actually computing such an integral. In standard, one-dimensional integration, we could use the fact that integrals are inverses of derivatives. We would like to use this idea here, too. To calculate a volume integral, we adapt the idea of a Riemann sum in one dimension. How we split V into volume elements δV should not matter so long as all volume elements are small. However, we can choose shapes for the δV's that make the calculation easier. The simplest and most useful way of splitting V into small bits is into cuboids[2] of side lengths $\delta x \times \delta y \times \delta z$ (figure 1),

$$\delta V = \delta x \delta y \delta z.$$

The mass M of the body is then the sum of $\rho \delta V$ over all the cuboids contained in the volume V (again, in the limit of small δx, δy and δz). There will be a small amount of discrepancy between the volume covered by cuboids and the actual volume V near the boundary of V, where we may either have cuboids that protrude slightly beyond V or parts of V that are not fully covered. This will necessarily be the case if the shape of V does not have flat surfaces. The discrepancy will however get smaller and smaller as we make size of the cuboids smaller.

[2]You may know a cuboid as a rectangular prism; the two mean the same thing.

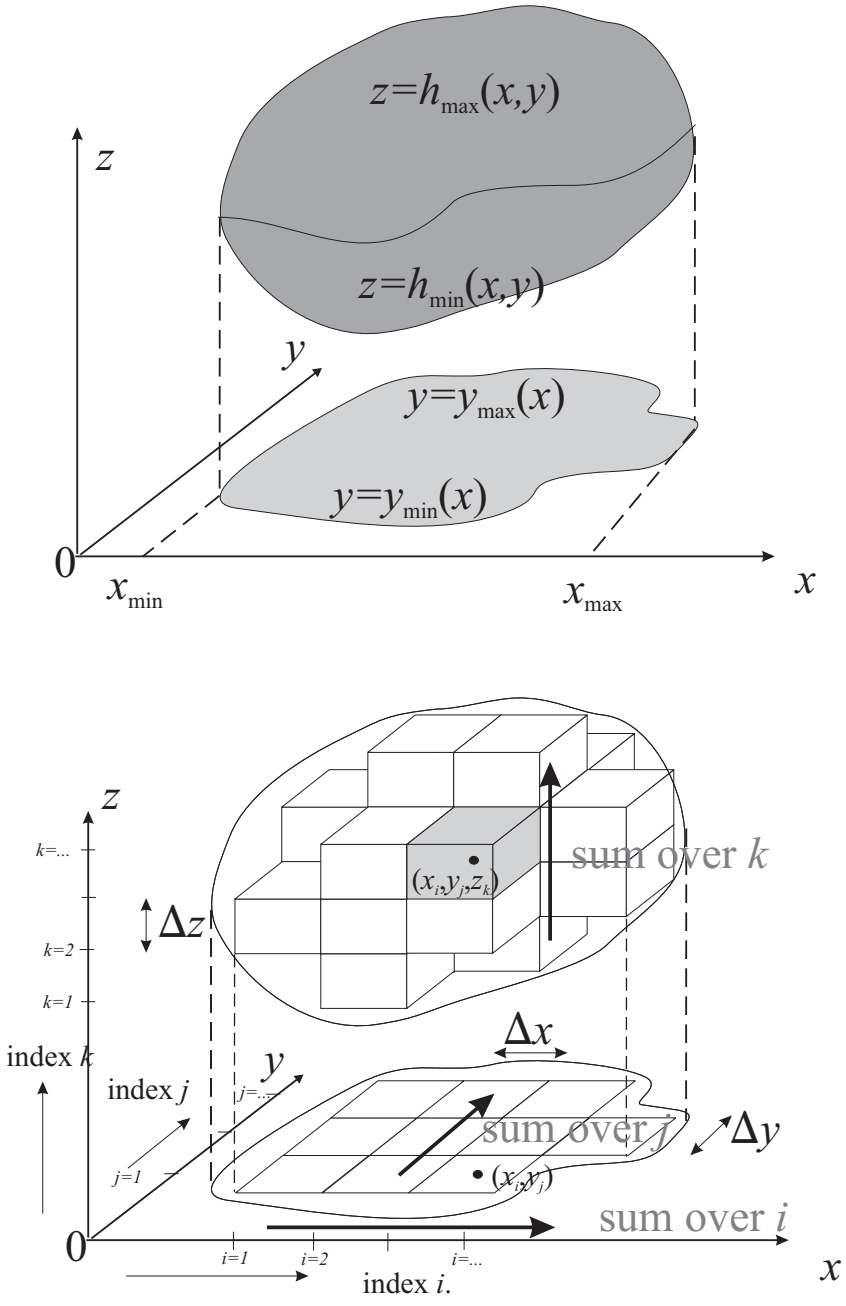

Figure 1: Breaking V into small cuboids δV.

To keep track of all the cuboids, we can use indices i, j and k to identify their x-, y- and z-positions respectively, letting x_i, y_j and z_k be the coordinates of the centre of the cuboid at the kth level above the base of the vertical column of cuboids with position i, j in the xy-plane.

To sum, we have to sum over all indices (i, j, k) that correspond to cuboids in the volume V,

$$\int_V \rho \, \mathrm{d}V = \sum_{i,j,k} \rho(x_i, y_j, z_k) \delta x \delta y \delta z.$$

When taking a sum over multiple indices, the order of summing does not matter. We can therefore sum over k first, then over j and finally over i,

$$\int_V \rho \, \mathrm{d}V = \sum_i \left\{ \sum_j \left[\sum_k \rho(x_i, y_j, z_k) \delta z \right] \delta y \right\} \delta x.$$

Of course, as $\delta z \to 0$, the sum over k is the sum over cuboids in the (i, j)th vertical column, where we can treat i and j temporarily as constants. This sum turns into an integral over z:

$$\sum_k \rho(x_i, y_j, z_k) \delta z \to \int \rho(x_i, y_j, z) \, \mathrm{d}z.$$

As we just stated, the sum is formed with i and j treated as constants, so the integral over z is take, with x_i and y_j treated as constants.

Next, the sum over j is then the sum over a row of such columns with fixed i, and turns into an integral over y taken at constant x_i, so

$$\sum_j \left[\sum_k \rho(x_i, y_j, z_k) \delta z \right] \delta y \to \int \left[\int \rho(x_i, y, z) \, \mathrm{d}z \right] \, \mathrm{d}y.$$

Here, the integral with respect to z is computed first, *treating y as well as x_i as constants*. Subsequently, the integral with respect to y is computed, treating x_i as constant.

Lastly the sum over i is finally the sum over all thee rows of vertical columns, and turns into an integral over x, so

$$\begin{aligned} \int \rho \, \mathrm{d}V &= \sum_i \left\{ \sum_j \left[\sum_k \rho(x_i, y_j, z_k) \delta z \right] \delta y \right\} \delta x \\ &\to \int \left\{ \int \left[\int \rho(x, y, z) \, \mathrm{d}z \right] \, \mathrm{d}y \right\} \, \mathrm{d}x \end{aligned}$$

As in the sum over indices i, j, k,, the integral with respect to z is computed first, *treating y as well as x as constants*, then the integral with respect to y is computed,

treating x as constant, and lastly, the integral with respect to x is calculated. For obvious reasons, the integral over x, y and z is also known as a *triple integral*.

Importantly, the volume V occupies a finite and prescribed portion of space, so the Riemann sums above turn into *definite integrals*: for each index (i, j and k) we sum only over those cubes δV that lie within the volume V (see figure 1). The hard part is therefore often to figure out what limits to use.

Example 1 *Let $\rho = 1+x^2+y^2+z^2$, and let V be the unit cube, for which $0 < x < 1$, $0 < y < 1$, $0 < z < 1$. Compute $\int_V \rho \, \mathrm{d}V$.*

In this case, the limits of integration are clear as each coordinate ranges from 0 to 1. Computing each integral in turn,

$$\begin{aligned}
\int_V \rho \, \mathrm{d}V &= \int_0^1 \left\{ \int_0^1 \left[\int_0^1 1 + x^2 + y^2 + z^2 \, \mathrm{d}z \right] \mathrm{d}y \right\} \mathrm{d}x \\
&= \int_0^1 \left\{ \int_0^1 \left[z + x^2 z + y^2 z + z^3/3 \right]_{z=0}^{z=1} \mathrm{d}y \right\} \mathrm{d}z \\
&= \int_0^1 \left\{ \int_0^1 4/3 + x^2 + y^2 \, \mathrm{d}y \right\} \mathrm{d}x \\
&= \int_0^1 \left[4y/3 + x^2 y + y^3/3 \right]_{y=0}^{y=1} \mathrm{d}x \\
&= \int_0^1 5/3 + x^2 \, \mathrm{d}x \\
&= 2
\end{aligned}$$

The tricky issue is therefore the question of limits of integration for irregularly-shaped bodies. We consider only a body that has a single upper surface of the form $z = h_{\max}(x, y)$, and a lower surface $z = h_{\min}(x, y)$ (figure 1). Let V_{xy} be the projection of this body onto the xy-plane, and assume that this projection has an upper boundary in the xy-plane of the form $y = y_{\max}(x)$, a lower boundary $y = y_{\min}(x)$, a left-hand edge at $x = x_{\min}$ and a right-hand edge at $x = x_{\max}$.

Note 2 *This way of specifying the shape of the body is the same as saying that a point (x, y, z) lies inside the body if $x_{min} < x < x_{max}$, $y_{min}(x) < y < y_{max}(x)$, $h_{min}(x, y) < z < h_{max}(z, y)$.*

There are more complicated bodies that cannot be put in this form. There may, for instance, be more than one upper and lower surface — think of a piece of swiss cheese with holes in the inside to picture this. If you need to deal with such a body, you can generally split that body (effectively, by cutting it) into simpler ones that conform to our requirements.

Note 3 *By 'projection', we mean the set of points (x, y) in the xy-plane that has at least some part of the body above it. In other words, a point (x, y) is in V_{xy} if there is*

a corresponding point (x, y, z) with the same x- and y-coordinates that lies inside the volume V. You can think of V_{xy} physically as the shadow the volume V would cast on the xy-plane when illuminated from far away by a very sharp light source positioned vertically above V.

The limits here come once more from viewing the integrals as sums. When taking the sum over k at fixed i and j, we are summing from the bottom of a vertical column to the top. This means that the range of k will depend on x_i and y_j, as z_k must lie between $h_{\min}(x_i, y_j)$ and $h_{\max}(x_i, y_j)$. In other words, when taking the limit $\delta z \to 0$, the integral is from $z = h_{\min}(x_i, y_j)$ to $h_{\max}(x_i, y_j)$. The next sum, over j, must capture all the columns in a row at fixed i, that is, the integral must range from $y_{\min}(x_i)$ to $y_{\max}(x_i)$. Lastly, the sum over i must go over all these rows, and the x-integral must therefore range from $x_{\min}$ to $x_{\max}$ (figure 1). Hence

$$\int_V \rho \, \mathrm{d}V = \int_{x_{\min}}^{x_{\max}} \left[\int_{y_{\min}(x)}^{y_{\max}(x)} \left(\int_{h_{\min}(x,y)}^{h_{\max}(x,y)} \rho(x, y, z) \, \mathrm{d}z \right) \mathrm{d}y \right] \mathrm{d}x. \tag{1}$$

We give an example of how to use this formula next, but before we do, it is worth re-iterating several points

1. When the volume integral is written in the form above, you have to do the 'innermost' integral in round brackets first. This is the integral with respect to z. When computing that integral, you have to treat x and y as constants

2. Next, you have to apply the limits to that integral. These are limits in z (meaning, you apply the limits $z = h_{\min}(x, y)$ and $z = \max(x, y)$ to the integral you have just done. Obviously, the limits can depend on x and y, just as ρ does. Again, you have to treat x and y as constants at this stage. If, after applying the limits, the expression you obtain still contains the variable z, you have made a mistake.

3. Completing the previous two steps should therefore have turned the integral in the round brackets into a function of x and y only. The next step is to compute the integral of that function with respect to y, treating x as constant. In other words, you can now compute the integral in the squared brackets, having figured out what the integrand is through the previous two steps.

4. Next, apply the limits, which can again depend on x, and you again treat x as a constant.

5. Taking the last two steps should turn the integral in square brackets into a function of x only. If the variable y still features, you have made a mistake. The next step is to compute the outermost integral with respect to x and apply the limits. The answer you get should not contain x, y or z

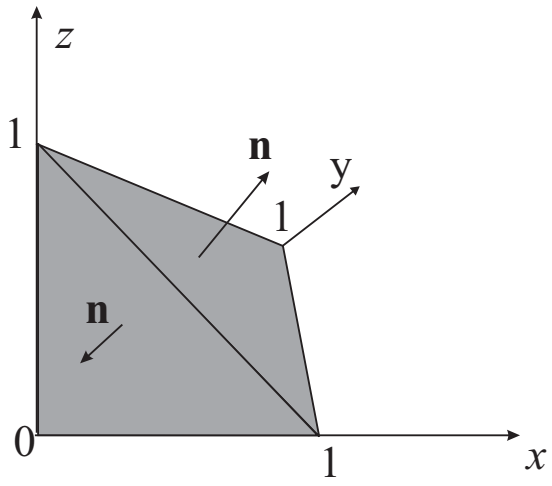

Figure 2: The volume V in example 2. (Ignore the vectors labelled $\mathbf{n}$ for now.)

Note 4 *There is of course nothing magic about the x-, y and z-directions, and the strict order in which you have to do the integrals above is dictated only by the way in which we have decided to specify the surface of the volume V. What we have done is to say that V consists of points (x, y, z) that satisfy the requirement that x lies between the left- and right-hand ends of the projection of V onto the xy-plane, $x_{\min} < x < x_{\max}$, y lies between the top and bottom boundaries of the projection y $y_{\min}(x) < y < y_{\max}(x)$, and z lies between an upper and a lower boundary, $h_{\min}(x, y) < z < h_{\max}(x, y)$.*

Suppose we wanted to work instead with the projection of V onto the yz plane and define the surface through a left-hand and right-hand boundary. In that case, we would have to say that y lies between the front and back ends of the projection, $\hat{y}_{\min} < y < \hat{y}_{\max}$, z lies between the top and bottom boundaries of the projection, $\hat{z}_{\min}(y) < z < \hat{z}_{\max}(y)$, and x between the left-hand and right-hand boundaries, $\hat{h}_{\min}(y, z) < x < \hat{h}_{\max}(y, z)$. The hats on $y_{\min}$, $y_{\max}$ etc. are just supposed to indicate that they are not the same as when we define V with the upper and lower boundary $h_{\min}(x, y)$ and $h_{\max}(x, y)$: for instance, the functions $h_{\min}(x, y)$ and $\hat{h}_{\min}(y, z)$ are not in general the same.

If we use the projection of V onto the yz plane and the left- and right-hand boundary of the volume as described above, write down the volume integral $\int_V \rho \, \mathrm{d}V$ as a triple integral analogous to (1). What is different? Which variable would you integrate with respect to first, second and third?

Example 2 *Let V be the tetrahedron with vertices $(0, 0, 0)$, $(1, 0, 0)$, $(0, 1, 0)$ and $(0, 0, 1)$ (figure 2), and let $\rho = x + y + z$. Compute $\int_V \rho \, \mathrm{d}V$.*

From the diagram, we have $h_{\min}(x, y) = 0$, $y_{\min}(x) = 0$, $x_{\min} = 0$, $x_{\max} = 1$. We still need to find $y_{\max}(x)$ and $h_{\max}(x, y)$.

Let us find $y_{\max}(x)$ first. Clearly, this is the straight line connecting the vertices $(0,1,0)$ to $(1,0,0)$. Let us write this in the form $y = ax+b$. If the line passes through $(0,1,0)$ then $y = 1$ when $x = 0$, so $1 = a \times 0 + b$, or $b = 1$. Similarly, $y = 0$ when $x = 1$ if the line passes through $(1,0,0)$, so $0 = a \times 1 + b$, or $a = -b = -1$. Hence

$$y_{\max} = 1 - x.$$

Next, we need to write the slanted upper surface in the form $z = h_{\max}(x,y)$. This works the same way as above. We know that the triangle is part of a plane, and hence we expect that

$$z = ax + by + c. \tag{2}$$

describes the surface for some set of coefficients a, b and c (with a and b potentially different from above). How do we find them? We know that the plane must pass through the points $(x,y,z) = (1,0,0)$, $(0,1,0)$ and $(0,0,1)$. Hence these coordinates must satisfy the equation (2). Substituting the coordinates into (2), we get, for each of these points in turn,

$$\begin{aligned} 0 &= a + c \\ 0 &= b + c \\ 1 &= c \end{aligned}$$

Hence $c = 1$, $a = b = -1$ and $z = 1 - x - y$. Therefore

$$h_{\max}(x,y) = 1 - x - y.$$

This allows us to write

$$\int_V \rho \, \mathrm{d}V = \int_0^1 \left[\int_0^{1-x} \left(\int_0^{1-x-y} (x+y+z) \, \mathrm{d}z \right) \mathrm{d}y \right] \mathrm{d}x.$$

But we have already stated that we will take the z-integral first, treating x and y as constant. This is done at follows:

$$\begin{aligned} \int_0^{1-x-y} (x+y+z) \, \mathrm{d}z &= \left[xz + yz + z^2/2 \right]_{z=0}^{z=1-x-y} \\ &= x(1-x-y) + y(1-x-y) + (1-x-y)^2/2 \end{aligned}$$

Note that when computing the integral with respect to z we treat x and y as constants not only in the integrand, but also in the limits. Once we have completely computed the integral with respect to z, including applying the limits*, we compute the integral with*

respect to y, treating x as constant. Last, the integral with respect to x is computed:

$$\begin{aligned}
\int_V \rho \,\mathrm{d}V &= \int_0^1 \left[\int_0^{1-x} \left(x(1-x-y) + y(1-x-y) + (1-x-y)^2/2 \right) \,\mathrm{d}y \right] \,\mathrm{d}x \\
&= \int_0^1 \left[x(1-x)y - xy^2/2 + (1-x)y^2/2 - y^3/3 - (1-x-y)^3/6 \right]_{y=0}^{y=1-x} \,\mathrm{d}x \\
&= \int_0^1 x(1-x) - x/2 + (1-x)/2 + (1-x)^3/6 \,\mathrm{d}x \\
&= \int_0^1 \left[1/2 - x^2 + (1-x)^3/6 \right] \,\mathrm{d}x \\
&= \left[x/2 - x^3/3 - (1-x)^4/24 \right]_0^1 \\
&= 1/8
\end{aligned}$$

Exercise 1 *Let V be a cuboid with rectangular base area A in the xy-plane and height h. Let density ρ be a function of height z only. Show that $\int_V \rho \,\mathrm{d}V = A \int_0^h \rho(z) \,\mathrm{d}z$.*

Exercise 2 *Let V be the tetrahedron with vertices $(0,0,0)$, $(1,0,0)$, $(1,1,0)$ and $(0,0,1)$. Let $\rho = 1 + xy + z$. Compute the mass of the body.*

Exercise 3 *What do you think the physical meaning of the integral $\int_V 1 \,\mathrm{d}V$ is? Confirm your answer by computing this integral for the tetrahedron V in the previous exercise, and compare that answer with the volume of the tetrahedron given by 1/3×base area×height.*

If $\rho = \rho_0 =$ constant, evaluate $\int_V \rho \,\mathrm{d}V$, if V also denotes the volume of the body V. (Comment?)

Density re-visited and generalized

You will have seen the analogy between density and velocity: both are constants of proportionality, one linking volume and mass, the other time and displacement. Both need to be defined *locally*: $\delta m = \rho \delta V$ and $\delta x = v \delta t$ only work when δV and δt are small.

For large volumes, we instead have

$$M = \int_V \rho \,\mathrm{d}V,$$

just as for large time intervals, we have

$$D = x(t_2) - x(t_1) = \int_{t_1}^{t_2} v(t) \,\mathrm{d}t.$$

If we take mass over volume, we get a mean density

$$\bar{\rho} = \frac{M}{V} = \frac{\int_V \rho \, \mathrm{d}V}{V},$$

but this is value depends on the particular shape of the volume V and does not give any information about the local density at a point (x, y, z). Similarly, the definition of velocity as distance over time gives a mean velocity — but not velocity at any instant in time:

$$\bar{v} = \frac{x(t_2) - x(t_1)}{t_2 - t_1}.$$

This can also be written in terms of an integral over the instantaneous velocity,

$$\bar{v} = \frac{\int_{t_1}^{t_2} v(t) \, \mathrm{d}t}{t_2 - t_1},$$

but really gives a mean velocity over the interval from t_1 to t_2

A better way to think of the density field $\rho(x, y, z)$ than as mass over volume is as a measure of how concentrated mass is near a point (x, y, z), which comes from the idea tha the mass δm of a small volume δV is related to that volume through $\delta m = \rho \delta V$. Similarly, velocity $v(t)$ is the rate of change of distance with respect to time, which comes from the idea that a displacement δx is related to time elapsed through $\delta x = v(t)\delta t$.

There are other quantities that are similar to density. For instance, the concentration of a chemical measured in moles per cubic metre, the heat content of a substance, measured in Joules per cubic metre, or charge density, measured in Coulombs per cubic metre. All of these quantities have in common that they measure some physical entity 'per unit volume', and their definitions make sense because, like mass, we expect chemical content, heat content or charge content to be proportional to volume, again provided the volume involved is small.

The theory developed above can also be applied to these other densities. For instance, we would define a chemical concentration field $c(x, y, z)$ in terms of the the number of moles δn in a small volume δV around the point (x, y, z),

$$c(x, y, z) = \frac{\delta n}{\delta V}.$$

Following the same steps as above, it is then clear that the number of moles N of the chemical in some fixed volume V is given by

$$N = \int_V c \, \mathrm{d}V.$$

Similarly, if $q(x, y, z)$ is the charge density field, the total charge Q in a given volume V is

$$Q = \int_V q \, \mathrm{d}V.$$

Because of this analogy, it makes sense to use the word 'density' to describe the concentration of some physical quantity in space, even if that quantity is not mass. In that case, depending on context, it may be worth referring to the original density as 'mass density'.

One important aspect of the densities we have mentioned above is that they are associated with conserved quantites, like mass, energy and charge. We need to know how conservation of these quantities constrains the way in which the densities can evolve in time or vary in space, and for this, we need to understand what a *surface integral* is.

Note 5 *Heat and amount of a particular chemical species are not by themselves conserved, but they also cannot be created or destroyed at will. To create heat — or* internal energy *to those familiar with thermodynamics, where 'heat' has a somewhat more specific meaning — requires another form of energy to be converted into heat. To create or destroy a chemical species requires the the destruction or creation of another chemical.*

Visualizing scalar fields

Density $\rho(x, y, z, t)$ is a 'field': a function of position and possibly of time. Density is also a scalar, so we call it a 'scalar field'. This can be contrasted for instance with velocity $\mathbf{v}$, which is a vector that may also depend on position and time. More on that later in the course.

To make sense of what a scalar field 'looks' like, we need a way of visualizing it. For an ordinary function of one variable, visualizing is relatively easy. If we have a function $y = f(x)$, we end up plotting y against x as a curve. This becomes harder when something is a function of two variables. $z = f(x, y)$ must be represented as a surface above the xy-plane. There is no extension of this to higher dimensions: we cannot plot $\rho(x, y, z)$ as the equivalent of a curve or surface.

A more common way to visualize a function even of two variables than plotting a surface — which is hard to do by hand —is to plot contours. If the function is given by $f(x, y)$, this means we plot a number of curves along which f is constant. In other words, we plot the sets of points (x, y) given by

$$f(x, y) = C_i$$

where the constant C_i is the ith contour level. Typically, a constant contour interval ΔC is used, so $C_2 - C_1 = C_3 - C_2 = C_i - C_{i-1} = \Delta C$ is constant.

Note 6 *A contour is also known as an* isoline *or as a* level set*. The concept can be extended to higher dimensions. We can define* isosurfaces *of a function of three variables. These would be the sets of points for which*

$$f(x, y, z) = C_i.$$

Sometimes, the easiest way to plot contours involves solving $f(x, y) = C_i$ for either x or y as a function of x, and using the sketching techniques we developed previously.

Example 3 *Let $f(x, y) = x^2 - y^2$. Suppose we want to plot contours of f with constant contour intervals, including the $f = 0$ contour as well as contours with negative and positive contour levels, if they exist.*

We are looking for contours defined through

$$x^2 - y^2 = C$$

The easiest thing to do seems to be to solve for y:

$$y = \pm\sqrt{x^2 - C}$$

It is essential that have the '$\pm$' sign here, otherwise we will get only have the contours. Of course, the $\pm$ means that each contour above the x-axis has a mirror image counterpart below x-axis, with the same contour level. We can therefore look at the $+$ case only, and fill in the mirror image contours afterwards. We therefore look only at

$$y = \sqrt{x^2 - C}.$$

What do these curves look like?

The easiest case is $C = 0$, in which case $y = \pm x$. We get two straight lines through the origin, inclined at $\pm 45^\circ$ to the coordinate axes. For $C \neq 0$, we can work through the list of steps needed to plot curves,

First, look at intercepts. The y-intercept is at $x = 0$, We can solve for y in that case only if $C < 0$, in which case we get $y = \sqrt{-C} = \sqrt{|C|}$. An x-intercept corresponds to $y = 0$, or $x = \sqrt{C}$ This exists only if $C > 0$. This suggests we should treat the $C < 0$ and $C > 0$ cases separately.

Let us look at $C < 0$ first, so negative contours. We do have a y-intercept at $\sqrt{C}$, and no x-intercept. There are no singularities, as y is defined for every x. To find asymptotes, note that as $x \to \infty$, x^2 becomes much larger than whatever the fixed contour level C is, so $\sqrt{x^2 - C} \approx \sqrt{x^2} = |x|$. The asymptotes are therefore $y = |x|$, regardless of the contour level C.

Next, we need to look at maxima or minima. We have

$$\frac{\mathrm{d}y}{\mathrm{d}x} = \frac{x}{\sqrt{x^2 - C}}.$$

This is zero when $x = 0$, so we have a stationary point there. This also happens to be the y-intercept.

We could also show that

$$\begin{aligned}\frac{\mathrm{d}^2y}{\mathrm{d}x^2} &= \left[\frac{1}{\sqrt{x^2-C}} - \frac{x^2}{(x^2-C)^{3/2}}\right] \\ &= \left[\frac{x^2-C}{(x^2-C)^{3/2}} - \frac{x^2}{(x^2-C)^{3/2}}\right] \\ &= \left[-\frac{C}{(x^2-C)^{3/2}}\right]\end{aligned}$$

with $C < 0$, this is positive and we have a local minimum. This would also have been obvious if we had looked at the fact that we have only one stationary point, and the cure asymptotes to $y = +|x|$.

We can now plot a single contour. We should note that the curve $y = \sqrt{x^2 - C}$ is symmetric about the y-axis, because x and $-x$ correspond to the same value of y.

As we change the contour level C, the interesting question is how the contour changes. Each will still have a minimum at the y-intercept $x = 0$ and be symmetric about the y-axis. Each contour also still asymptotes to the same straight line $y = |x|$, which happens to be the $C = 0$ contour. The only thing that changes is that the y-intercept moves up as we change C, with $y = \sqrt{|C|}$. The y-intercept therefore does not change proportionally to C; for small C, the jump in y-intercept between adjacent contours is larger than for large C. The picture that emerges is show in figure 3.

We still have to deal with $C > 0$, and eventually fill in the mirror image contours below the x-axis. The latter is trivial to do. To do the former, we could go through the entire procedure above again. It is however simpler to use symmetry again. The function $f(x, y) = x^2 - y^2$ is antisymmetric *under changes of x and y: For instance $f(2, 1) = 2^2 - 1^2 = 3$ while $f(1, 2) = 1^2 - 2^2 = -3 = -f(2, 1)$. More generally $f(y, x) = -f(x, y)$. Note that the point (y, x) is the mirror image of the point (x, y) about the line $y = x$ (which passes through the origin and lies at 45°) This means that, if a point (x, y) lies on a contour for some value of C, then its mirror image (y, x) lies on the $-C$ contour. If we know contours above that line, we can fill in the contours below that line simply by reflection. The same is true for reflection about the line $y = -x$: we also have $f(y, x) - f(-x, -y)$.*

The full contour plot therefore looks as in figure 4.

Exercise 4 *Sketch constant interval contours for the following functions, including the $f = 0$ contour. Be sure to include contours for f negative* and *positive, if they exist.*

1. $f(x, y) = x^2 + y$
2. $f(x, y) = x^2 + 2xy - y^2$
3. $f(x, y) = xy + x^2$

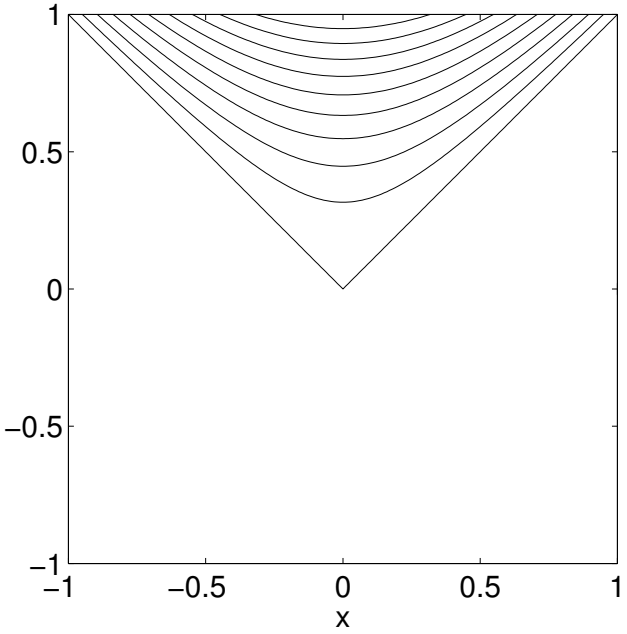

Figure 3: The contours given by $y = +\sqrt{x^2 - C}$ with $C \leq 0$. These contours have *negative* (or zero) contour levels.

In some cases, it is easier not to solve for y as a function of x or x as a function of y. This is often the case when there is for instance rotational symmetry, as is the case when functions $f(x, y)$ depend on x and y only through $r = \sqrt{x^2 + y^2}$. r is the distance from a point (x, y) to the origin, and if f depends only on r, we immediately know that contours are circles. The only question that arises then is what the contour spacing looks like.

Exercise 5 *Sletch constant interval contours for the following functions*

1. $f(x, y) = x^2 + y^2$
2. $f(x, y) = \exp(-x^2 - y^2)$

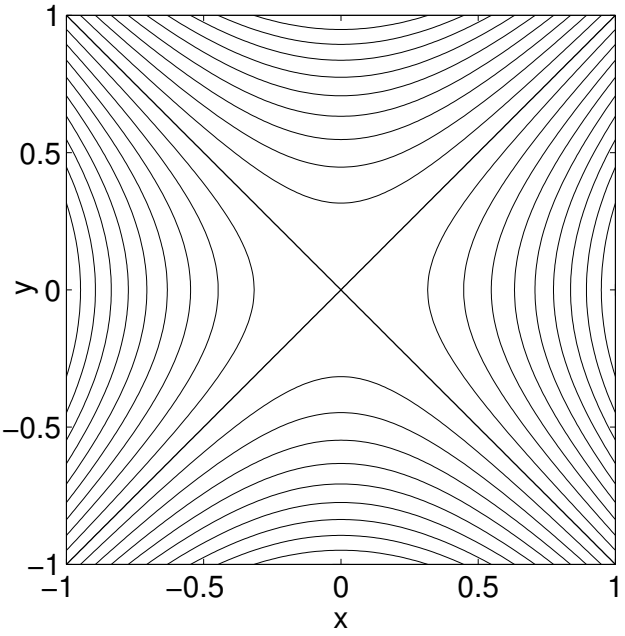

Figure 4: Contours of $f(x, y) = x^2 - y^2$. Positive values of f are found along the x-axis, negative values along the y-axis. If this were an elevation map, the origin would be a saddle, valleys would lie to the 'north' and 'south' with mountain ridges to the 'east' and 'west'.

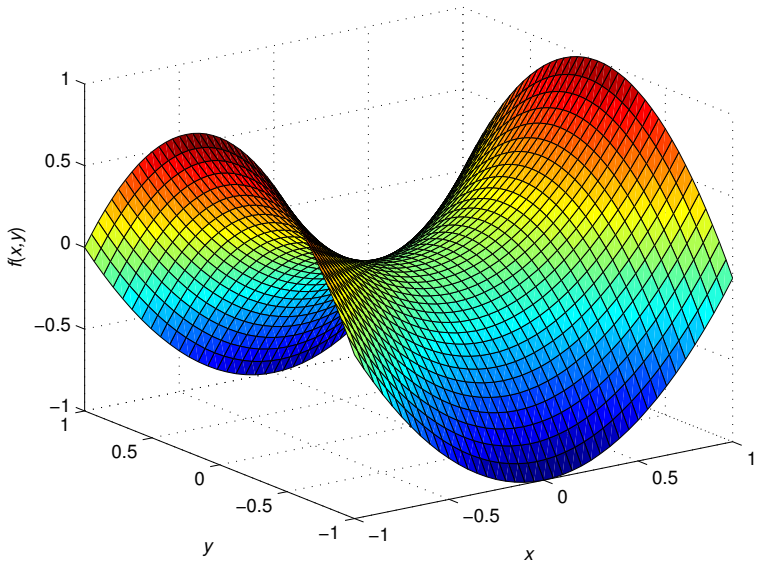

Figure 5: As we are dealing with a function of two variables, we can also plot the surface given by $f(x, y) = x^2 - y^2$. This is however much harder to do by hand.