What is Integration Good For?

Or how I learned to stop worrying and love the integral,

Or why the heck are we learning all of this?

Marty Weissman
April 19, 2000
Math 21a

This handout is to explain what an integral "means." This handout does not explain how to evaluate integrals. This handout will hopefully help you convert concepts into mathematical language. This handout will not help you use that mathematical language to find answers.

Part I: Finding the size of a region

Single variable calculus: Let $y = f(x)$ be a function of one variable. You should be familiar with integrals of the form: $\int_a^b f(x)dx$. If the function $f(x)$ is positive, then this (by definition) computes the area between the x-axis and the graph of $f(x)$ over the ***domain of integration*** $a \leq x \leq b$. Remember, every integral involves a domain of integration, and an "integrand," i.e. the thing being integrated.

We can also use integrals (inefficiently in this case) to find the size of the domain of integration. Here in the one-dimensional case, if the domain of integration is given by $a \leq x \leq b$, then the size of the domain of integration is just the length of the interval $[a,b]$, which is just $b-a$. We could alternatively compute the size of the domain of integration by evaluating the integral:

$$\int_a^b 1 \cdot dx = x\Big|_a^b = b - a .$$

Two-variable calculus: Let $z = f(x,y)$ be a function of two variables, and let R be a "filled-in" region in the plane. You should be familiar with integrals of the form: $\iint_R f(x,y)dA$. We will call integrals of this kind ***abstract area integrals***, to emphasize the fact that we don't specify whether $dA = dxdy$ or whether $dA = rdrdq$, and we don't give a parametrization of the region R in terms of coordinates. If the function $f(x,y)$ is positive, then this integral evaluates the volume of the solid contained above the region R in the xy-plane, and below the graph of $z = f(x,y)$.

We can also use integrals (sometimes efficiently in this case) to find the size of the domain of integration, i.e. the area of the region R. Just like the one-variable case, we have the formula:

$$\iint_R 1 \cdot dA = Area(R).$$

To evaluate an abstract double integral, you have to convert it into an iterated integral by choosing coordinates (Cartesian or polar), expressing the region R in terms of these coordinates, substituting $x = r\cos(\theta), y = r\sin(\theta)$ in the case of polar coordinates, and expressing dA in the proper coordinates.

Three-variable calculus: Let $f(x, y, z)$ be a function of three variables, and let R be a solid region in space. You should be somewhat familiar with integrals of the form: $\iiint_R f(x, y, z)dV$. Such an expression is called an ***abstract volume integral***, to emphasize the fact that we don't specify whether $dV = dxdydz$ or $dV = rdrd\theta dz$ or $dV = r^2\sin(\varphi)d\theta d\varphi dr$, i.e. we don't yet make a choice of coordinates, nor do we choose a parametrization of the solid region R. The reason you might be somewhat familiar with these integrals, rather than being completely familiar, is that they don't compute a volume that is easy to visualize. They do compute the volume of a four-dimensional region, between the "xyz-hyperplane" and the graph of $w = f(x, y, z)$, but we (unfortunately) don't worry about that in our class.

In the special case when $f(x, y, z) = 1$ however, this sort of integral should certainly be familiar. Integrating 1, as usual, computes the size of the domain of integration. Here, we have:

$$\iiint_R 1 \cdot dV = Volume(R).$$

Arc-length-integrals: Let γ be a curve in the plane (or in space), parametrized by $\vec{x}(t) = (x(t), y(t))$ (or $\vec{x}(t) = (x(t), y(t), z(t))$). Define the ***length differential*** by $d\ell = \left|d\vec{x}\right| = \sqrt{x'(t)^2 + y'(t)^2}dt$ (or $\sqrt{x'(t)^2 + y'(t)^2 + z'(t)^2}dt$). Let $f(x, y)$ be a function of two variables (or $f(x, y, z)$ of three variables). Then you should be somewhat comfortable with integrals of the form: $\int_\gamma f(x, y)d\ell$ (or $\int_\gamma f(x, y, z)d\ell$), called **abstract arc-length integrals**. They are abstract, since there is no parametrization of γ written down in the integral, nor an expansion of $d\ell$ in terms of a parameter. Such an integral computes the area of the "curtain", hanging down from the graph of $z = f(x, y)$ to the curve γ in the xy-plane.

In the special case when we integrate 1, which might be the case you're more used to from this course, the integral again computes the size of the region. In other words, if the curve γ doesn't overlap itself, then:

$$\int_\gamma 1 \cdot d\ell = length(\gamma).$$

Surface-area-integrals: Let M be a surface lying in space, e.g. the surface of a sphere, a plane, the surface of a donut, etc… Let $f(x, y, z)$ be a function of three variables. Then you should be getting comfortable with integrals of the form $\iint_M f(x, y, z)dS$, called ***abstract surface-area integrals***. They are abstract, since there is no parametrization of M such as $\vec{X}(u, v) = (x(u, v), y(u, v), z(u, v))$ written down in the integral, nor do we write $dS = \left|\overrightarrow{X}_u \times \overrightarrow{X}_v\right| dudv$. Once again, such an integral does not compute a useful area or volume, unless you're capable of imagining three-dimensional curtains hanging in four-dimensional space.

In the special case when we integrate 1, this sort of integral will compute the size of the domain of integration: M. That is to say:

$$\iint_M 1 \cdot dS = Area(M).$$

Conclusion: Single, double, and triple integrals can be used to find the size of a one, two, or three dimensional region by integrating 1. Areas and volumes beneath graphs of functions can be computed by integrating the functions. More generally, areas and volumes between graphs of two functions can be computed by integrating the difference of the two functions. Arc-length integrals, and surface integrals are useful for computing arc-length and surface area. There are often many ways of finding the size of a region – only practice will develop the instinct for which ones are easier.

Part II: Densities and Total "Masses"

Very often in applications, a region will have a certain variable density. A one-dimensional region (for instance a wire) might have a charge density which tells us how many electrons there are per foot of wire. Moreover, this density may vary from one position on the wire to another. A surface, for instance the surface of the Earth, might have a pollutant density, which tells us how many grams of pollutants there are per square meter on the surface of the Earth. This sort of density will also often vary from region to region. Finally, a region in space, such as a galaxy, will have a mass-density (i.e. the usual density), which describes how many kilograms of matter are contained in a cubic parsec. This density will vary widely from the center of the galaxy, to black holes, to empty tracts of space with only a few specks of dust.

Integration allows us to pass from a density to a total. Integration allows us to go from knowing the charge density on a wire to the total number of electrons on the wire. Integration allows us to go from the pollutant density on the Earth to how much total crap there is. Integration allows us to go from measurements of the density of the galaxy to the total mass of the galaxy.

Example II.1: An electrode is a metal disc, with radius 1 inch. Hooking up the electrode to a power supply charges the disc with a varying charge density of $\rho = e^{-r}$ coulombs per square inch, where r is the distance from a point to the center of the disc. What is the total charge on the electrode?

Answer: The total charge is the integral of the density: $\iint_R \rho dA$, where R is the unit disc in the plane, and ρ is given as above.

Example II.2: The starship Enterprise is hanging out in the gamma-quadrant, when all of a sudden everything starts shaking, the red-alert goes off, and Captain Picard is sent flying into the turbo-lift. Data turns to you and says, "We've just run across a cosmic filament, given by a curve g , parametrized by the equation $\vec{x}(t) = (t, \sin(t), \cos(t))$, where $0 \le t \le 1729$, with a density given by $r(t) = 1 + t$ megagrams per meter." What is the mass of this cosmic filament?

Answer: The total mass is the integral of the density: $\int_g r\, d\ell$, with g , r as above.

Example II.3: A galaxy (say Andromeda) is approximately a cylinder, of radius 10000 parsecs, and height 30 parsecs. The (mass) density of the galaxy at a point at a distance r from the axis of the galaxy, and at height z (varying between –15 and 15) is given by $r(r,z) = h^2(e^{-r} + r^2)$ megagrams per cubic parsec. What is the mass of the galaxy?

Answer: The total mass is the integral of the density: $\iiint_G r\, dV$, where G is the cylindrical region described above.

Part III: Averages

Averages are important concepts. They don't require much introduction. Chances are, if you're going to use mathematics for future applications, averages will be there in some form or another. When the sample is discrete, e.g. a bunch of individuals, you form an average by adding up values and dividing by the sample size. When the sample is continuous, e.g. all the locations on the surface of the Earth, you form an average by integrating and dividing by the sample size, e.g. the surface area of the Earth.

Example III.1: The temperature on the surface of the Earth at a given location is approximately $T(q, j) = 30\cos(j)$ degrees Celsius, where q, j are the usual spherical coordinates. What is the average temperature on the Earth?

Answer: The average temperature is given by the quotient of integrals:

$$Ave(T) = \iint_E T(q, j)\, dS \Big/ Area(E)$$

Here E denotes the surface of the Earth.

Example III.2: A pond is given by a cylindrical region full of water, of radius 100 feet, and depth 13 feet. If r, d represent the distance from a point in the pond to the center of the pond, and the depth of a point respectively, then the pH at a given point is given by $A(t) = 7 + e^{-r^2}$. What is the average pH of the pond?

Part IV: Flows and Fluxes

Flows and fluxes are extremely important applications of integration. Flows and fluxes always involve a vector field, i.e. a vector at each point in the plane or in space that represents force, or the current of a river, or wind-speed, or magnetic field, etc… We can compute the flow of a field along any curve, and we can compute the flux of any field through a curve in the plane or a surface in space.

Flows: Let $\vec{F}(x, y)$ be a vector field in the plane. We could easily consider vector fields and flows in space just by adding another variable in everything to come. Let g be a path in the plane, parametrized by $\vec{x}(t) = (x(t), y(t))$. The ***flow*** of $\vec{F}$ along g is defined to be the line-integral: $\int_g \vec{F} \cdot d\vec{x}$. If we travel along the path g, then this integral tells us how much we are going with the flow or against it. If the integral gives a large negative number, then we're doing a lot of work to move against the flow, i.e. it's tough to travel this path. If the integral gives a positive number, then the current does more work than we do – we're going with the flow. The specific number is a measure of work; it is work in the physics sense of the word if $\vec{F}$ denotes force in the physics sense of the word.

Example IV.1: An airplane flies from St. Louis to Chicago with position vector given by $\vec{x}(t) = (70t, 240t, \sin(pt)$, where $0 \le t \le 1$ is in hours, and distance is measured in miles. The wind speed is given by the vector field $\vec{F}(x, y, z) = (\cos(x) + z, \sin(y), 0)$. The airplane will consume fuel at a rate $r = 500 - \vec{v} \cdot \vec{F}$ gallons per hour, where $\vec{v}$ denotes the velocity at a given time. What is the total fuel consumed?

Answer: The total fuel consumed is given by the flow-type integral: $500 - \int_g \vec{F} \cdot d\vec{x}$.

Fluxes: Let $\vec{F}(x, y, z)$ be a vector field in space (here we'll restrict ourselves to three variables, since we used two in the last section). Let M be a surface in space, parametrized by $\vec{X}(u, v)$. The ***flux*** of $\vec{F}$ through M is a measure of how much the vector field $\vec{F}$ passes through M. It is computed via a surface integral:

$$\iint_M \vec{F} \cdot \vec{n} dS$$

Here $\vec{n}$ is a unit normal vector to the surface, and we can compute this using the formula: $\vec{n} \cdot dS = (\vec{X}_u \times \vec{X}_v) dudv$. Note that this is a cross product, and **not** its magnitude. If we imaging $\vec{F}$ as the current in a river, and M as a net in the water, then the flux will be the amount of water flowing through the net per unit time. Fluxes could be used to measure how much electricity flows through an objects, how much radiation passes through a person's skin in an x-ray machine, how much heat is released from the surface of a sun, etc… Also, fluxes are used in the flux-capacitor, which was used for time-travel in 1985.

Example IV.2: A carbon rod is dunked in radioactive water in a nuclear power plant in order to keep the reactions under control. The rod is a cylinder, of radius 1 foot, and height 20 feet, assumed to have the unit circle in the xy-plane as its bottom. Radioactive particles are flying around in the water, according to the vector field

$\vec{F}(x,y,z) = (-y, 2x, z)$, with units in feet per second. How many particles are entering the cylinder through its sides (not the top and bottom) per second?

Answer: The number of particles entering per second is given by the flux integral: $\iint_M \vec{F} \cdot \vec{n} dS$, where M is the lateral surface of the cylinder.

Glossary

Abstract Integrals: Often called just definite integrals in the literature, I like to distinguish between integrals with and without specified parametrizations. An abstract integral (in my language) is an integral *without* specified parametrizations. In other words neither the domain of integration, nor the differentials are parametrized.

Domain of integration: The region over which an integral is taken. A 1-dimensional domain of integration is a line segment or more general curve in the plane or in space. A two dimensional domain of integration is a portion of the xy-plane, or a portion of a surface in space. A three-dimensional domain of integration is a solid region in space.

Parametrization: The assignment of coordinate(s) (the parameter(s)) to a region. A parametrization of a curve is a representation of that curve in terms of a coordinate (usually *t*), for instance parametrizing the unit circle by the equation: $\vec{x}(t) = (\cos(t), \sin(t))$ Parametrizing a domain of integration means replacing the domain by the parameter limits. Parametrizing a differential means replacing symbols such as $dA, dV, d\ell$, by coordinate differentials such as dx, dy, dz, dt, du, dv in an appropriate way.