MATH2000

CALCULUS AND LINEAR ALGEBRA II

Lecture Workbook

Semester 1 2013

Contents

1 Solutions of first order ODEs

By the end of this section, you should be able to answer the following questions about first order ODEs:

- How do you solve an IVP associated with directly integrable, separable or linear ODEs? (Revision)
- Under what conditions does a solution to an IVP problem exist?
- Under what conditions is a solution to an IVP problem unique?

In MATH1052, you were introduced to Ordinary Differential Equations (ODEs) and Initial Value Problems (IVPs) and saw how to find solutions to some special types of first order equations. In particular, there should be three types of first order ODEs that you are familiar with solving.

- Directly integrable: $\dfrac{dy}{dx} = f(x)$.

- Separable: $\dfrac{dy}{dx} = f(x)g(y)$.

- Linear: $\dfrac{dy}{dx} = q(x) - p(x)y$.

In most applications involving first order ODEs, we are required to solve an IVP. Generally, this is a problem of the form

$$\frac{dy}{dx} = f(x, y), \quad y(x_0) = y_0.$$

In otherwords, we seek to find solutions of the ODE which pass through the point (x_0, y_0) in the x-y plane.

Consider the following three examples.

1.1 Example: $\frac{dy}{dx} = x$, $y(0) = 1$ has a unique solution

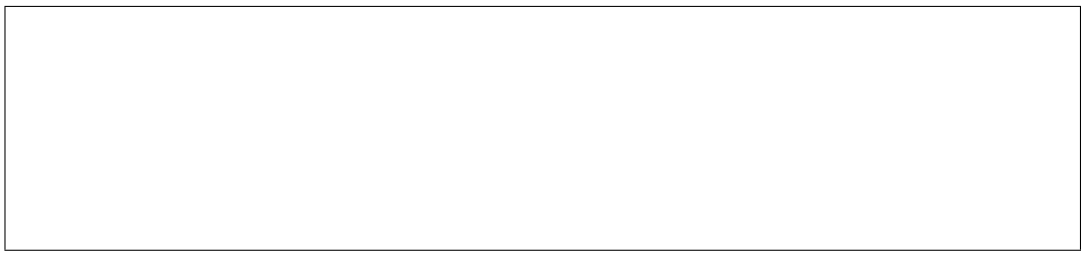

1.2 Example: $\frac{dy}{dx} = 3xy^{1/3}$, $y(0) = 0$ has more than one solution

1.3 Example: $\frac{dy}{dx} = \frac{x-y}{x}$, $y(0) = 1$ has no solution

1.4 Existence and uniqueness criteria

Here we consider the initial value problem

$$\frac{dy}{dx} = f(x,y), \quad y(x_0) = y_0.$$

- (existence) If $f(x,y)$ is continuous in some rectangle

  $$R = \{(x,y)|\ |x - x_0| < a, |y - y_0| < b\}$$

  then the initial value problem has at least one solution. Note that $(x_0, y_0) \in R$.

- (uniqueness) Moreover, if $f_y(x,y)$ is also continuous in R then there is at most one solution to the initial value problem.

The above two conditions only tell us that a solution exists or is unique locally (i.e., in the rectangle R). Beyond R, we simply don't know. Let's look at the previous three examples in the context of the theorem.

1.5 Example: $\frac{dy}{dx} = x$, $y(0) = 1$

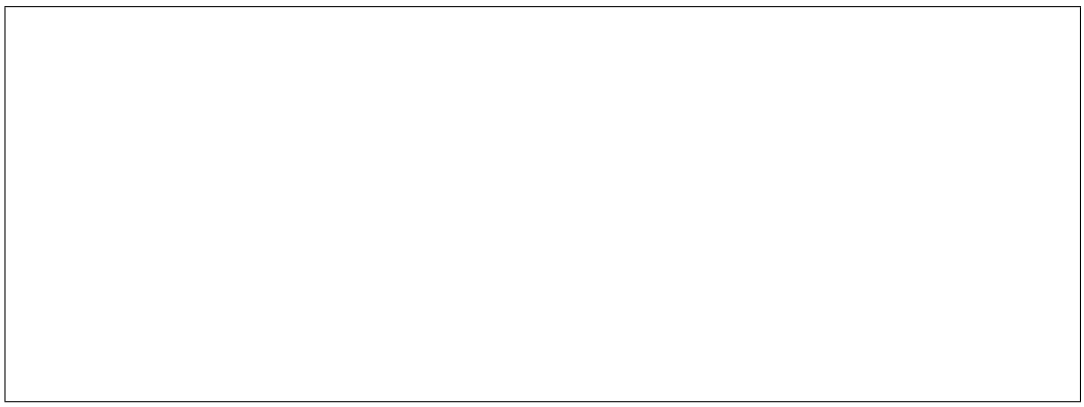

1.6 Example: $\frac{dy}{dx} = 3xy^{1/3}$, $y(0) = 0$

1.7 Example: $\frac{dy}{dx} = \frac{x-y}{x}$, $y(0) = 1$

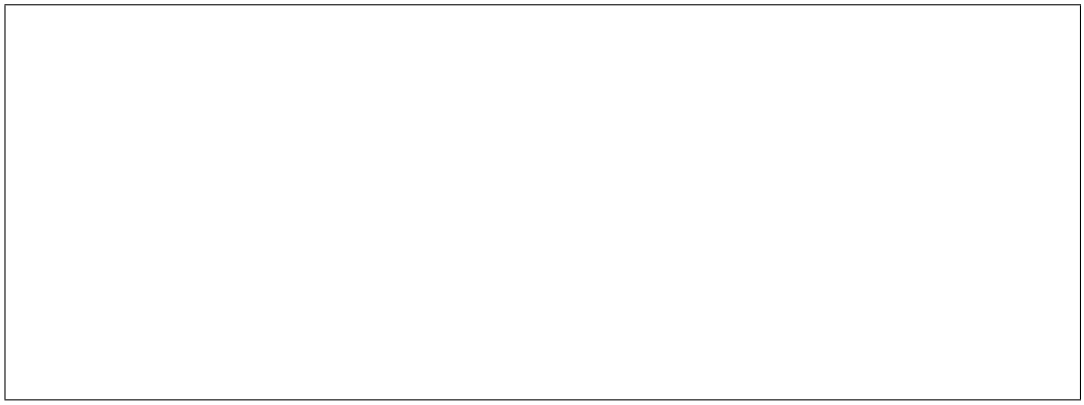

2 Exact first order ODEs

By the end of this section, you should be able to answer the following questions about first order ODEs:

- How do you identify an exact ODE?
- How do you solve an exact ODE?

2.1 Definition

First recall that if $z = f(x, y)$ is a differentiable function of x and y, where $x = g(t)$ and $y = h(t)$ are both differentiable functions of t, then z is a differentiable function of t whose derivative is given by the chain rule:

$$\frac{dz}{dt} = \frac{\partial f}{\partial x}\frac{dx}{dt} + \frac{\partial f}{\partial y}\frac{dy}{dt}.$$

Now suppose the equation

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

defines y implicitly as a function of x (here C is a constant). Then $y = y(x)$ can be shown to satisfy a first order ODE obtained by using the chain rule above. In this case, $z = f(x, y(x)) = C$, so

$$\begin{aligned} (0 = \ ) \quad \frac{dz}{dx} &= \frac{\partial f}{\partial x}\frac{dx}{dx} + \frac{\partial f}{\partial y}\frac{dy}{dx} \\ \Rightarrow \quad f_x + f_y y' &= 0. \end{aligned} \tag{1}$$

A first order ODE of the form

$$P(x, y) + Q(x, y)\frac{dy}{dx} = 0 \tag{2}$$

is called exact if there is a function $f(x, y)$ (compare (2) with (1) above) such that

$$f_x(x, y) = P(x, y) \text{ and } f_y(x, y) = Q(x, y).$$

The solution is then given implicitly by the equation

$$f(x, y) = C.$$

The constant C can usually be determined by some kind of "initial condition".

Given an equation of the form (2), how do we determine whether or not it is exact? There is a simple test.

2.2 Test for exactness

Let P, Q, $\dfrac{\partial P}{\partial y}$, and $\dfrac{\partial Q}{\partial x}$ be continuous over some region of interest. Then

$$P(x,y) + Q(x,y)\frac{dy}{dx} = 0$$

is an exact ODE iff

$$\frac{\partial P}{\partial y} = \frac{\partial Q}{\partial x}$$

everywhere in the region.

The problem of actually determining $f(x,y)$ is still outstanding. Consider the following example.

2.3 Example: $2x + e^y + xe^y y' = 0$

2.4 Almost exact ODEs and integrating factors

Let's say that we have an equation

$$P(x,y) + Q(x,y)\frac{dy}{dx} = 0$$

such that

$$\frac{\partial P}{\partial y} \neq \frac{\partial Q}{\partial x}.$$

The test we have just seen tells us that the ODE is not exact. Are we still able to do anything with it? Here we consider using an "integrating factor", which is different to the one introduced to solve linear ODEs.

The idea is to multiply the ODE by a function $h(x,y)$ and then see if it is possible to choose $h(x,y)$ such that the resulting equation

$$h(x,y)P(x,y) + h(x,y)Q(x,y)\frac{dy}{dx} = 0$$

is exact. We know from the test that this new equation is exact if and only if

$$\frac{\partial}{\partial y}(hP) = \frac{\partial}{\partial x}(hQ).$$

Let's see if we can find such a function:

In general, the equation for $h(x,y)$ is usually just as difficult to solve as the original ODE. In some cases, however, we may be able to find an integrating factor which is a function of only one of the variables x or y. Let's try $h \equiv h(x)$:

2.5 Example: $(3xy + y^2) + (x^2 + xy)\frac{dy}{dx} = 0$

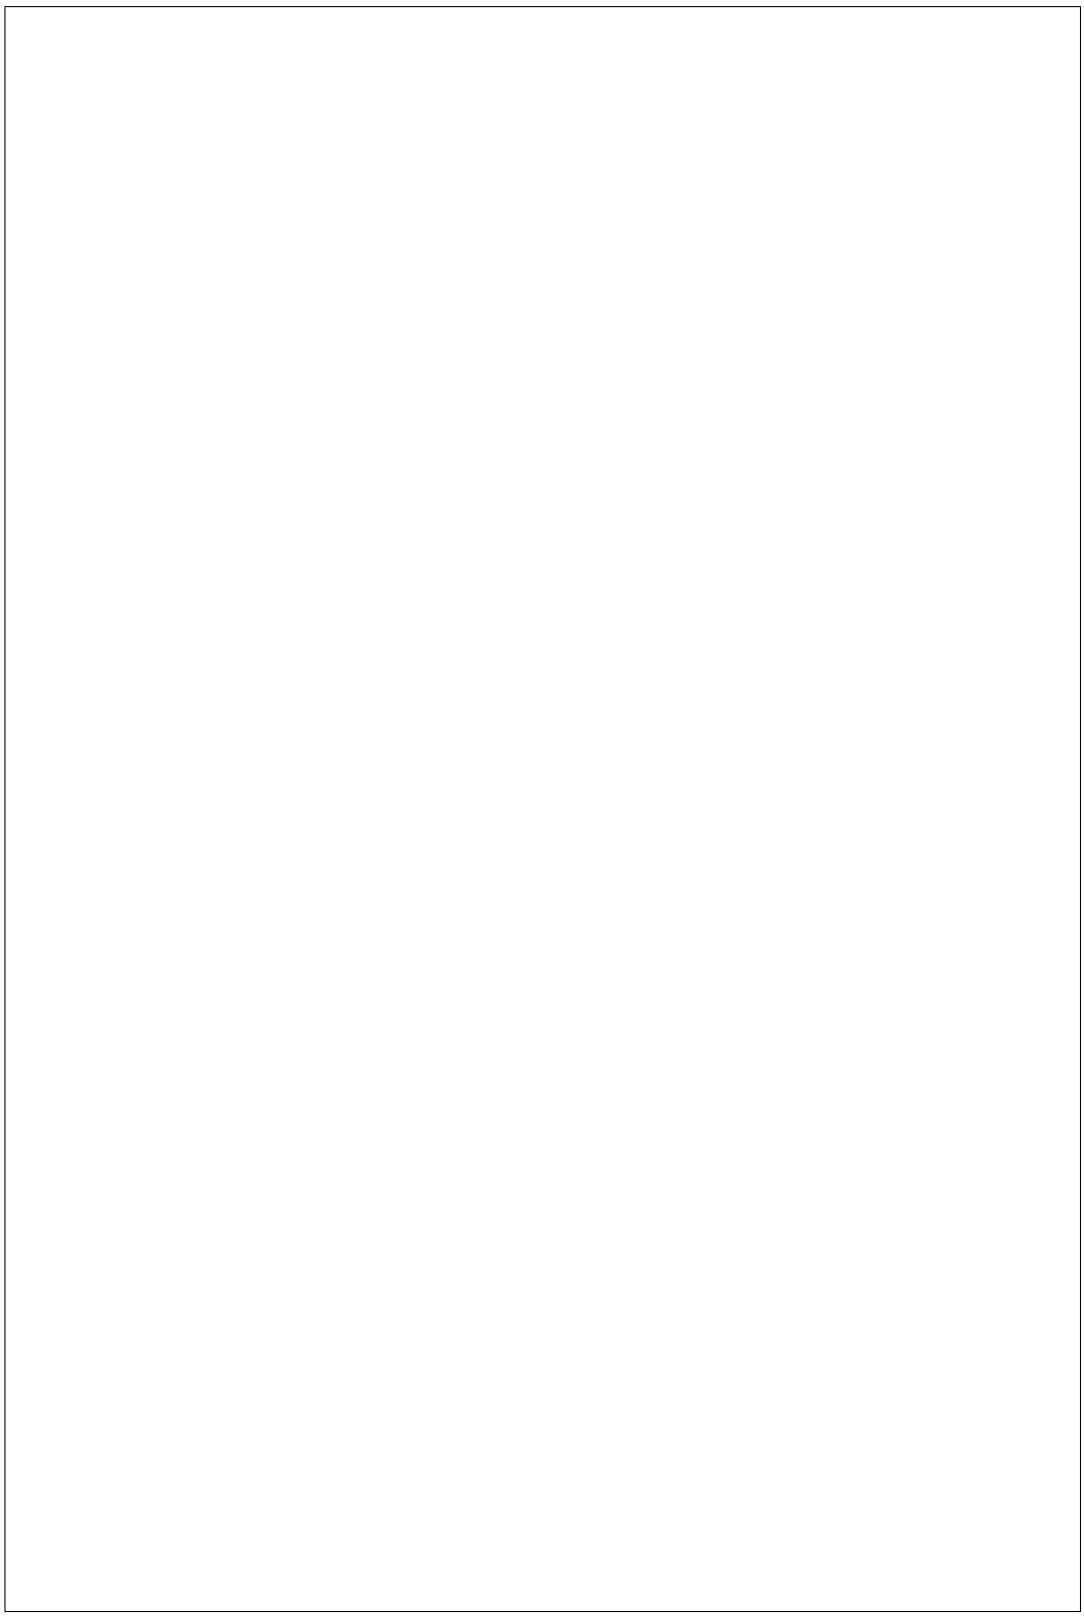

3 Linear second order nonhomogeneous ODEs, method of undetermined coefficients

By the end of this section, you should be able to answer the following questions:

- How do you apply the method of undetermined coefficients to solve a nonhomogeneous linear second order ODE?
- Under what conditions will the method work?

ODEs can be split into two classes: linear and non-linear. Non-linear ODEs are generally very difficult to solve. Linear ODEs are simpler because their solutions have general properties which facilitate working with them. There are also well established methods for solving many linear ODEs of practical significance.

A second order ODE is called linear if it can be written in the form

$$y'' + p(x)y' + q(x)y = r(x). \tag{3}$$

Any second order ODE which cannot be written in this form is called non-linear. Note that y and its derivatives appear linearly and p, q and r can be any functions.

Over the next few sections we study linear second order ODEs. The motivation for studying second order ODEs is twofold. Firstly they have applications in mechanics and electric circuit theory, so anyone studying either of these fields will most likely come across second order ODEs. Secondly, the theory of linear second order ODEs is very similar to that of higher order linear ODEs, so that the transition to studying higher order linear ODEs would not require too many new ideas.

Second order linear ODEs were introduced in MATH1052, and here we first recall some important results.

3.1 The superposition principle

If $r(x) = 0$ in equation (3), then we call the equation homogeneous. If $r(x) \neq 0$, the ODE is nonhomogeneous.

For any homogeneous linear equation, if y_1 and y_2 are solutions, so too is the linear combination $Ay_1 + By_2$. This is called the superposition principle. It is important to note that the superposition principle is not true for nonlinear equations and nonhomogeneous.

3.2 General solutions and initial value problems (homogeneous)

The general solution of a homogeneous linear ODE

$$y'' + p(x)y' + q(x)y = 0$$

is of the form

$$y = Ay_1 + By_2,$$

that is, a linear combination of two linearly independent solutions with two arbitrary constants A and B.

An initial value problem consists of a homogeneous (in this case) linear second order ODE and two initial conditions

$$y(x_0) = K_0, \quad y'(x_0) = K_1.$$

3.3 Homogeneous ODEs with constant coefficients

Let a, b be constants. We look at solving the ODE

$$y'' + ay' + by = 0. \tag{4}$$

By assuming the solution is of the form $y = e^{\lambda x}$, we conclude that λ satisfies the quadratic

$$\lambda^2 + a\lambda + b = 0.$$

This quadratic is called the characteristic equation (or auxillary equation) of (4), the roots of which are given by

$$\lambda_\pm = \frac{-a \pm \sqrt{a^2 - 4b}}{2}.$$

The form of the general solution depends on the roots of the characteristic equation, summarised in the table below.

Roots	General Solution
real distinct λ_+, λ_-	$y = Ae^{\lambda_+ x} + Be^{\lambda_- x}$
single real $\lambda = \alpha$	$y = (A + Bx)e^{\alpha x}$
complex $\lambda_\pm = \beta \pm i\omega$	$y = e^{\beta x} A\cos\omega x + e^{\beta x} B\sin\omega x$

3.4 Method of undetermined coefficients

Now we consider equations of the form

$$y'' + p(x)y' + q(x)y = r(x), \quad r(x) \neq 0. \tag{5}$$

You should know from MATH1052 that the general solution on an open interval I is

$$y = y_H + y_P,$$

where y_H is the general solution of the homogeneous equation (with $r(x) = 0$) on I and y_P is a particular solution of (5) on I containing no arbitrary constants.

In what follows, we determine a solution to the homogenous equation, and then *try* a form (with *undetermined coefficients*) for the particular solution which looks like it will result in the function on the right hand side.

The method of undetermined coefficients, as presented here, only works for the constant coefficient case:

$$y'' + ay' + by = r(x),$$

and $r(x)$ contains exponentials, polynomials, sines and cosines, or sums and certain products of these functions.

Choose for y_P a form similar to $r(x)$, involving unknown coefficients. The coefficients are then determined by substituting y_P into the ODE.

$r_i(x)$	$g_i(x)$	$r_i(x)$	$g_i(x)$
$ke^{\gamma x}$	$ae^{\gamma x}$	$k\cos\omega x,$ $k\sin\omega x$	$a\cos\omega x + b\sin\omega x$
$\sum_{i=0}^{N} k_i x^i,\ N = 0, 1, 2, \ldots$	$\sum_{i=0}^{N} a_i x^i$	$ke^{\alpha x}\cos\omega x,$ $ke^{\alpha x}\sin\omega x$	$e^{\alpha x}(a\cos\omega x + b\sin\omega x)$

3.5 Rules for method of undetermined coefficients

We follow these basic steps.

1. Find a solution y_H to the corresponding homogeneous equation.
2. For $r(x) = r_1(x) + r_2(x) + \ldots + r_n(x)$, we first make a guess $g(x) = g_1(x) + g_2(x) + \ldots + g_n(x)$ for y_P, where the $g_i(x)$ correspond to the $r_i(x)$ entries in the table above.
3. If a term $g_i(x)$ appears in y_H, replace $g_i(x)$ in the initial guess by $xg_i(x)$.
4. If any of the $xg_i(x)$ from step 3 appear in y_H, replace $xg_i(x)$ by $x^2 g_i(x)$.
5. Substitute the modified guess $g(x)$ into the left hand side of the ODE and equate coefficients on both sides. Once you have worked out the coefficients, the guess $g(x)$ becomes y_P.

3.6 Example: $y'' + 4y' + 4y = 8x^2$

The method looks relatively simple, but there are a number of well known special cases which the rules deal with. Consider the following two examples:

3.7 Example: $y'' + y' - 2y = -3e^{-2x}$

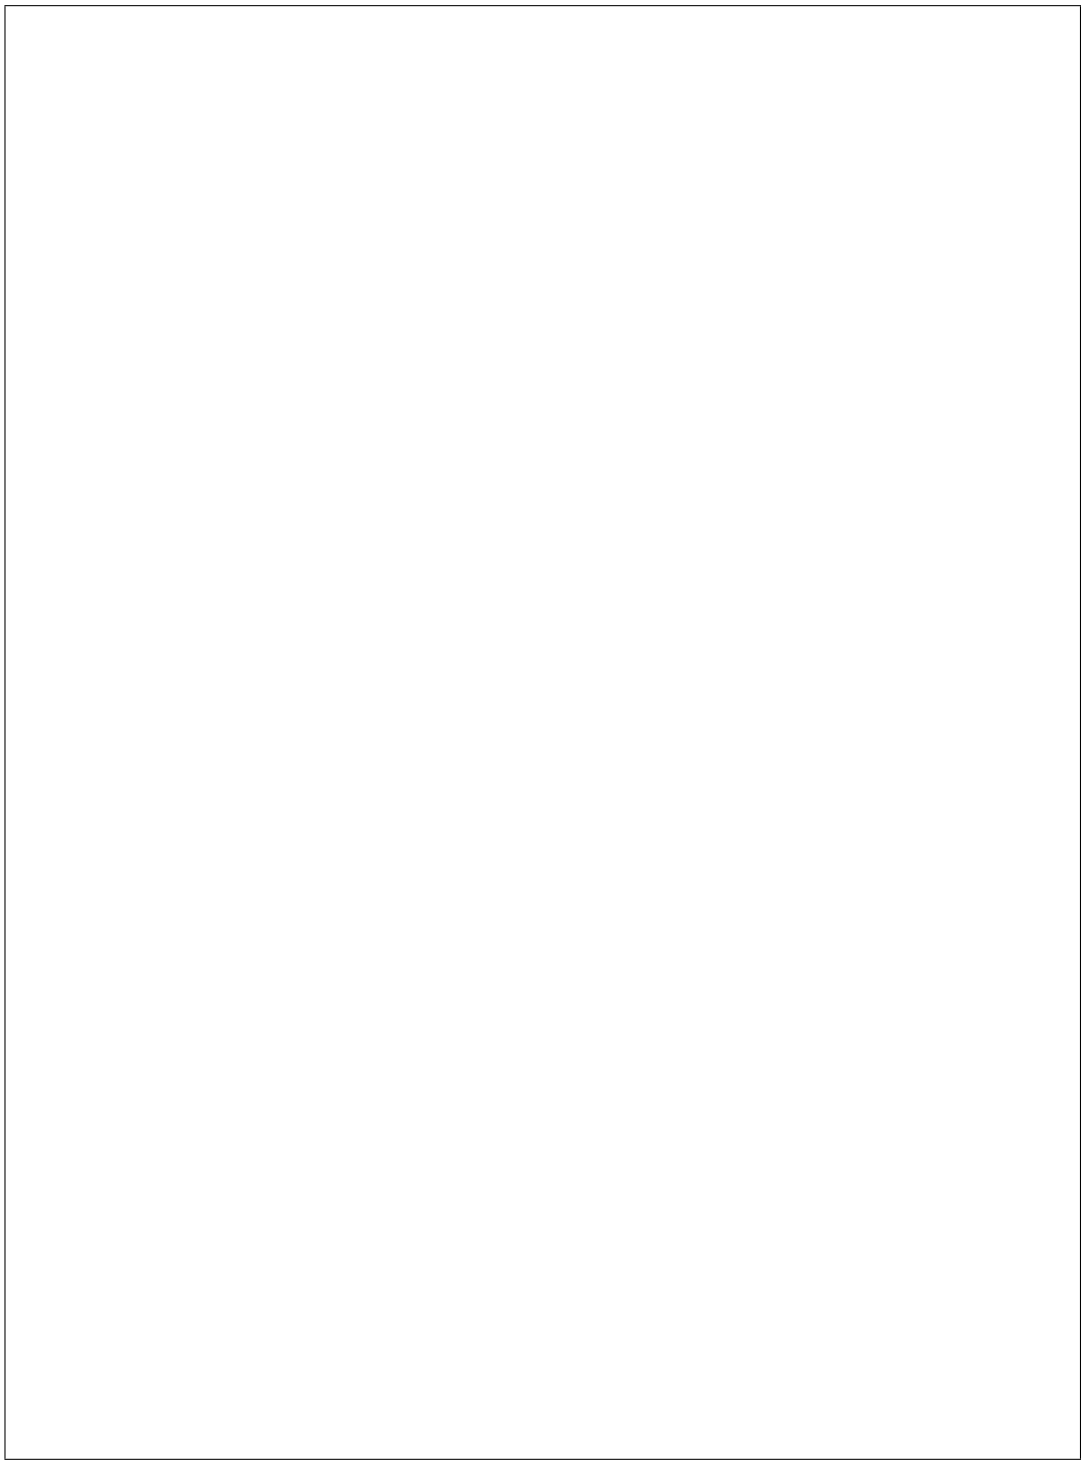

3.8 Example: $y'' - 2y' + y = e^x$

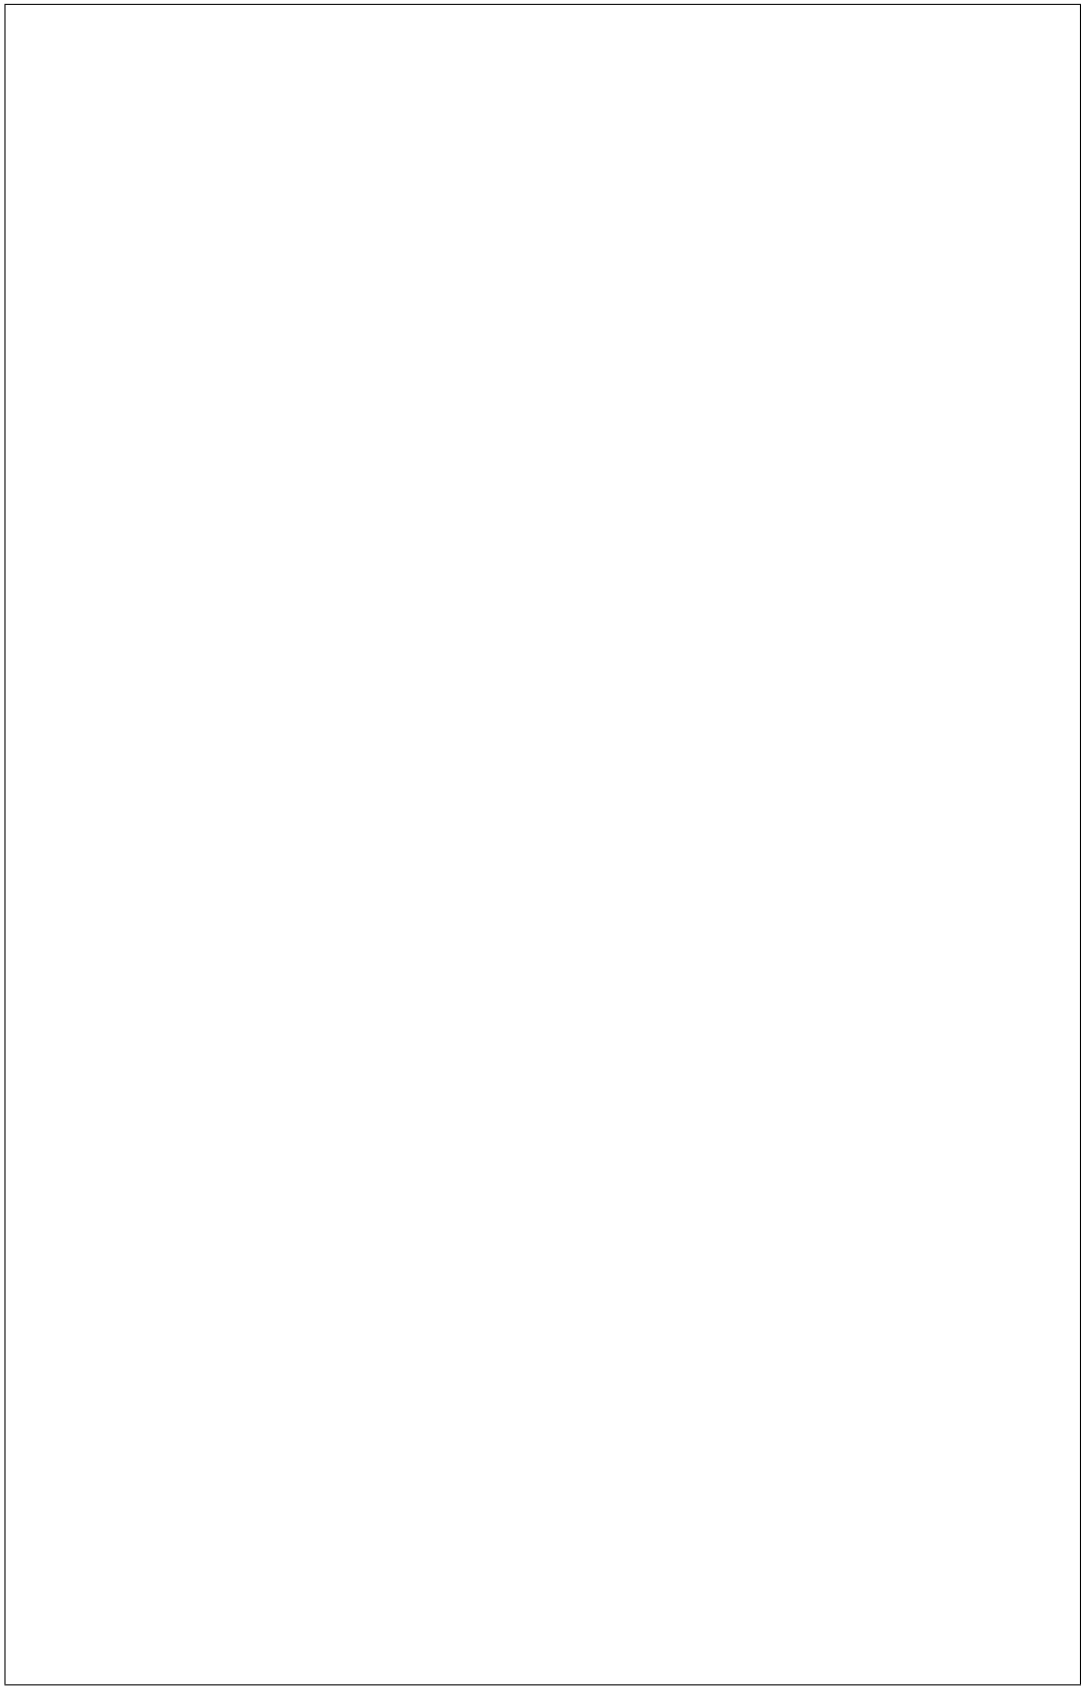

Note that if there are two terms on the right hand side we can handle each term separately. Consider the following extension of the previous example.

3.9 Extended example: $y'' - 2y' + y = e^x + x$

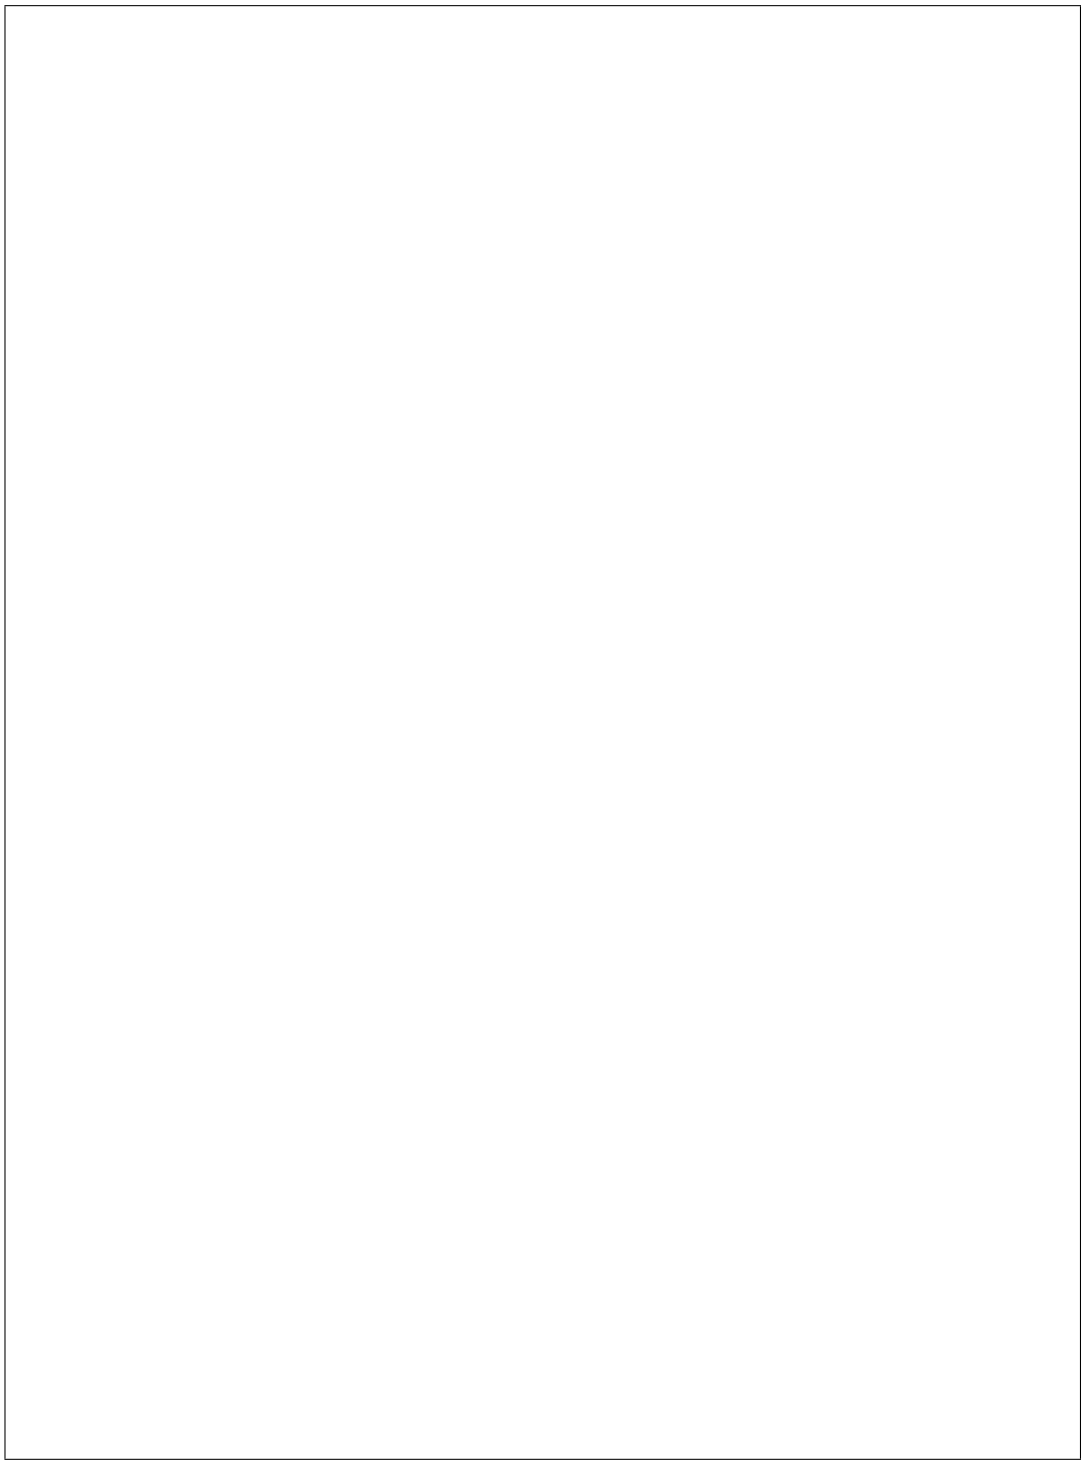

4 Variation of parameters

By the end of this section, you should be able to answer the following questions:

- Under what conditions does the method work?
- What functions need to be determined first before using the method?
- How do you use the variation of parameters method to solve a nonhomogeneous linear second order ODE?

The method of undetermined coefficients is very easy to apply, but only works for constant coefficients with certain $r(x)$. In the case

$$y'' + p(x)y' + q(x)y = r(x)$$

has arbitrary coefficient functions p, q, r, the variation of parameters works all the time. The process is the following:

- Solve $y'' + p(x)y' + q(x)y = 0$ to obtain a basis of solutions y_1, y_2 and set $W = y_1y_2' - y_1'y_2$ (this quantity is known as the Wronskian of the solutions y_1 and y_2). There is a result that states that $W \neq 0$ if and only if y_1 and y_2 are linearly independent.
- Set $y_P = u(x)y_1(x) + v(x)y_2(x)$ and substitute into the ODE. We also impose the condition $u'y_1 + v'y_2 = 0$. We have the freedom to impose this extra arbitrary condition because we have two functions (u and v) and only one equation they need to satisfy arising from the ODE.
- We obtain
$$u(x) = -\int \frac{y_2 r}{W} dx, \quad v(x) = \int \frac{y_1 r}{W} dx.$$

This approach is a variant of the method of Reduction of Order, which prescribes that we take a solution, say y_1 of the associated homogeneous equation and seek a particular solution of the form $y_p = U(x)y1$.

4.1 Derive the formulae for $u(x)$ and $v(x)$ in the variation of parameters

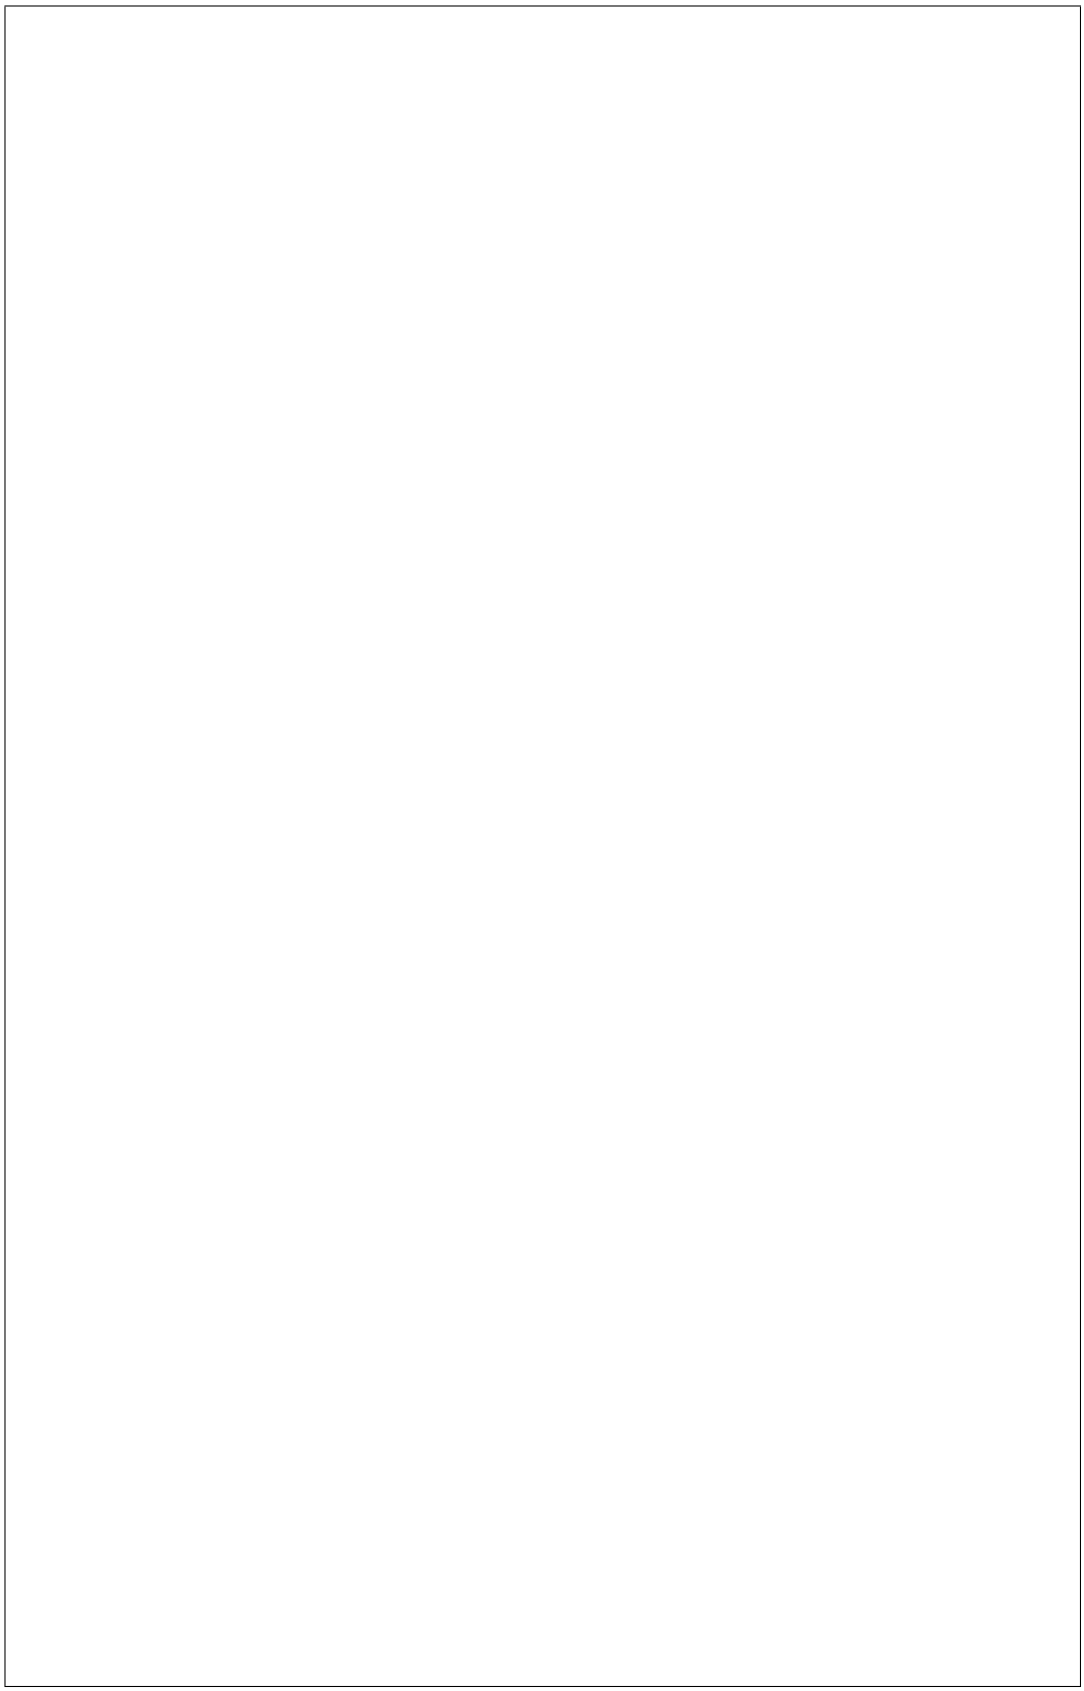

4.2 Example: $y'' - 4y' + 5y = 2e^{2x} / \sin x$

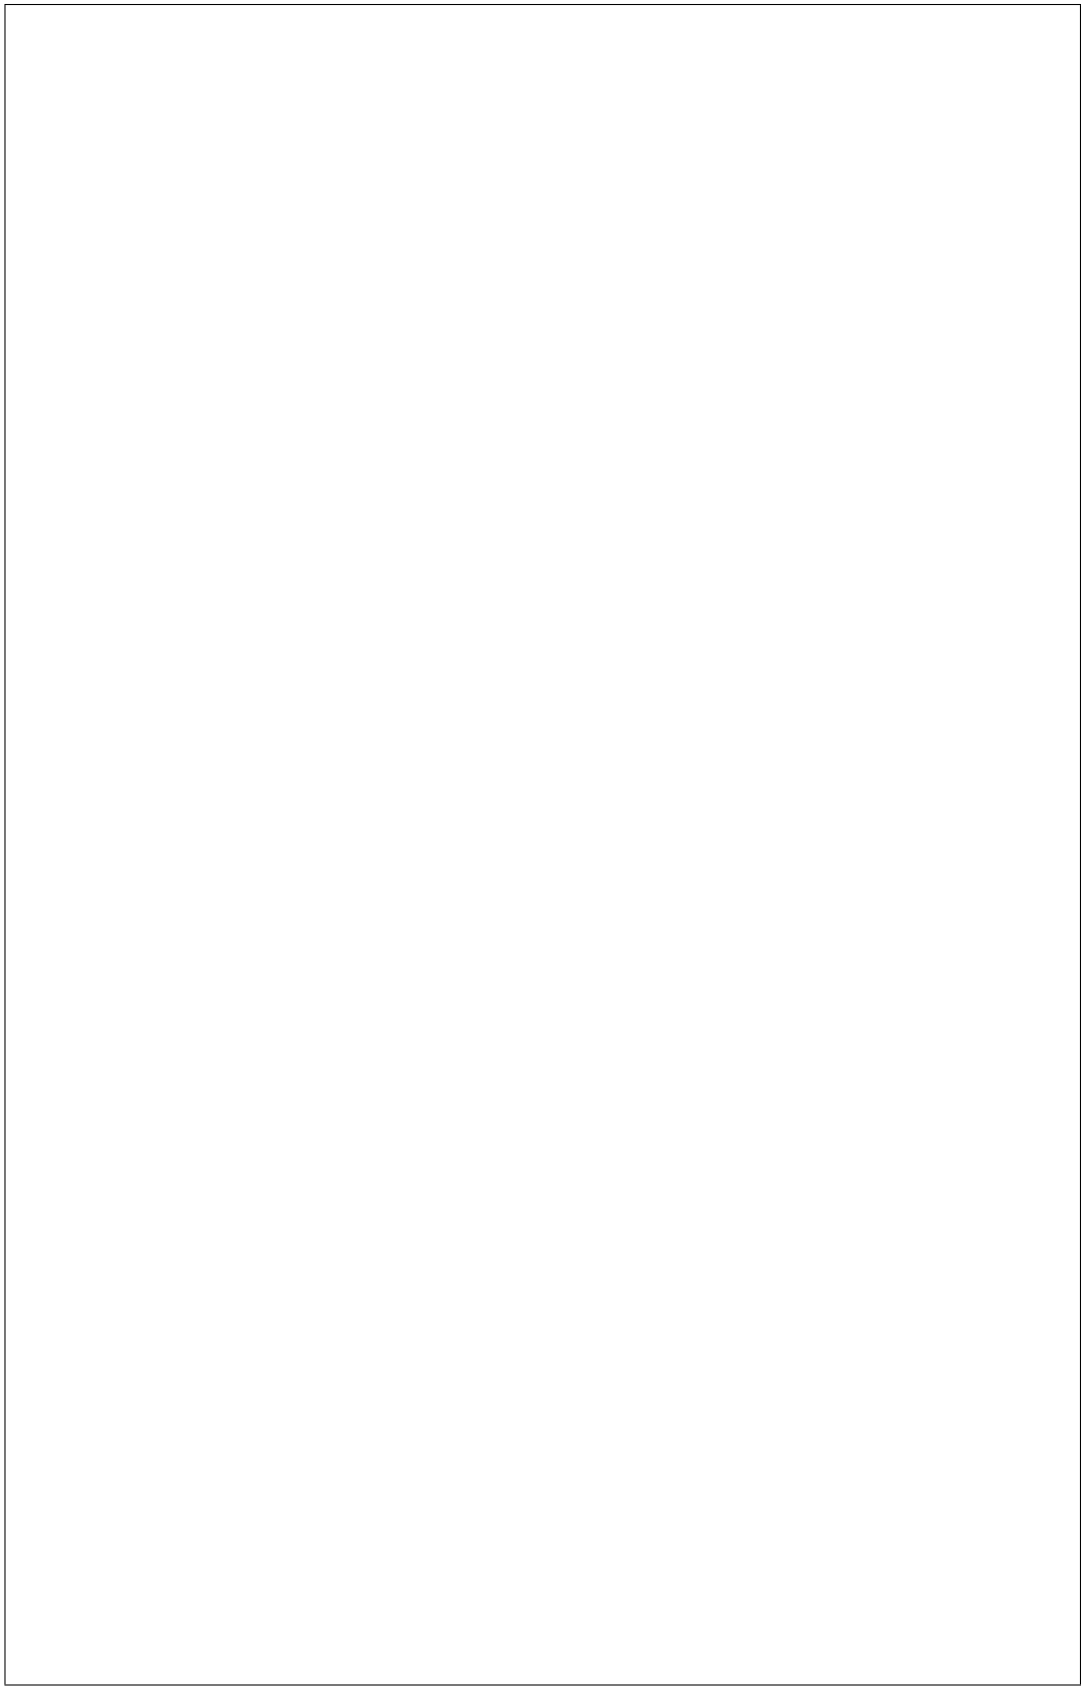

4.3 Example: $y'' + 4y = \csc 2x$

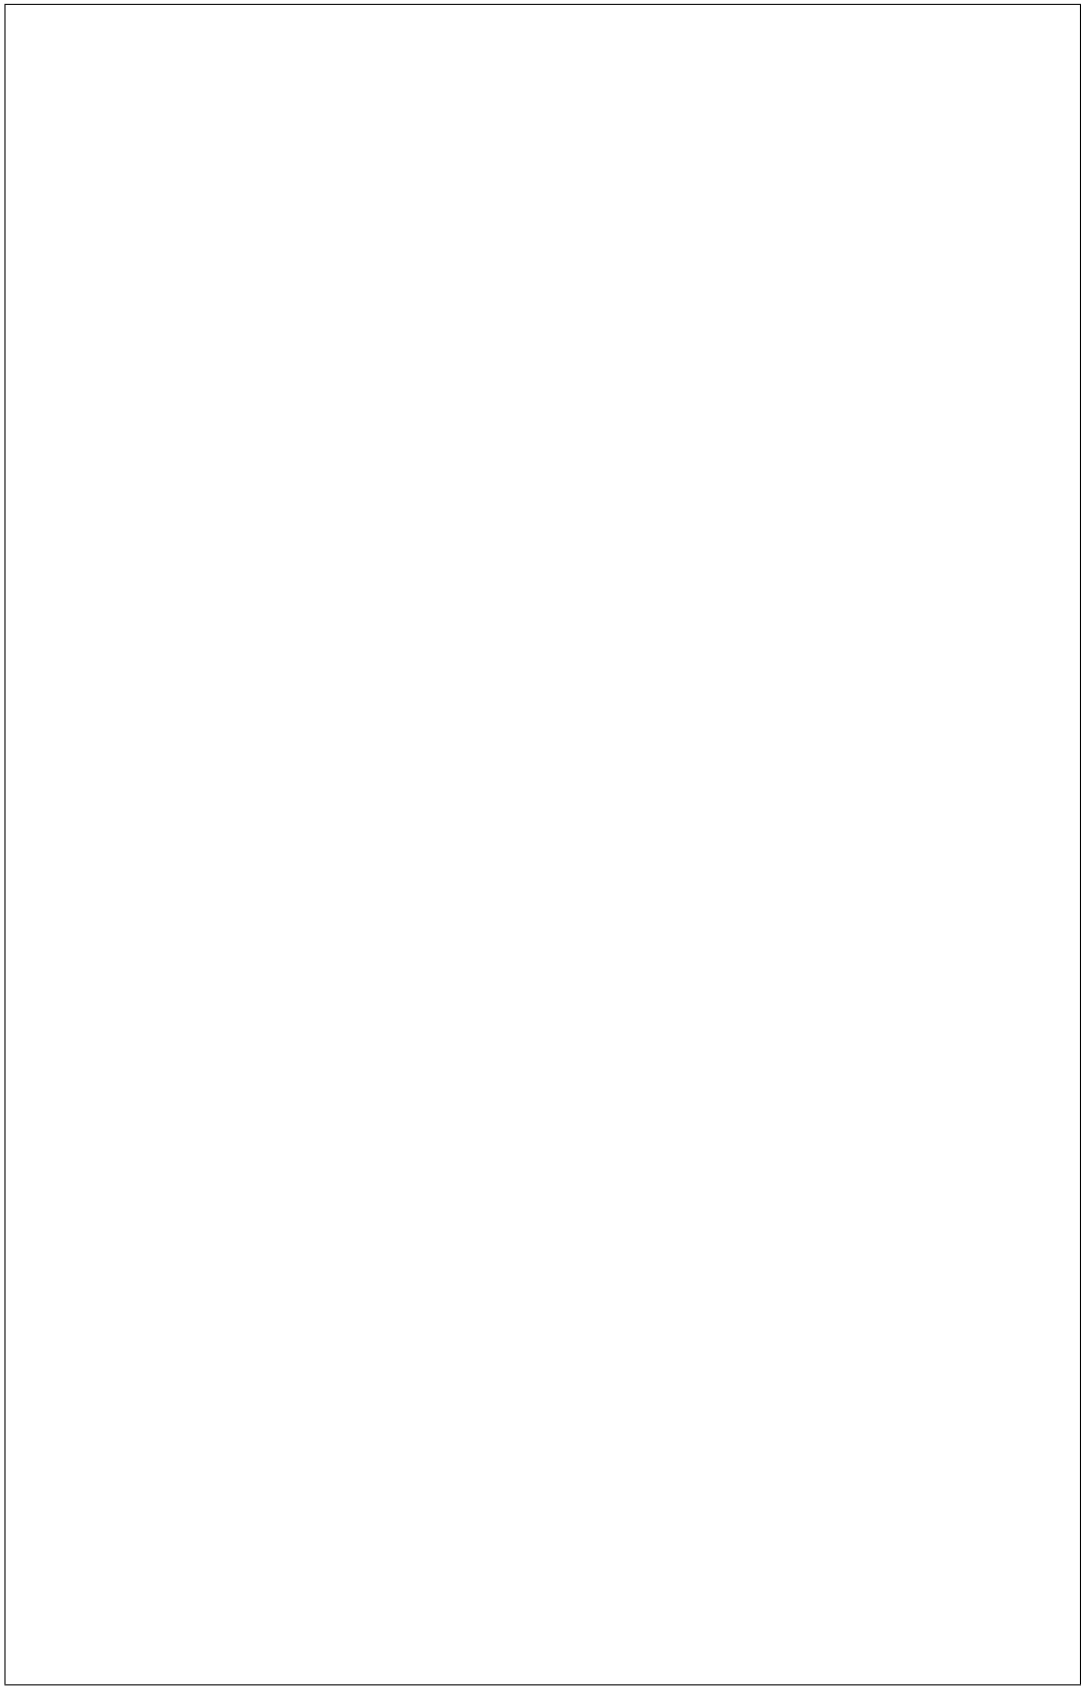

4.4 Summary of ODE techniques and types of equations you should know

- First order, directly integrable
- First order, separable
- First order, linear, integrating factor
- First order existence and uniqueness criteria
- First order, exact
- Second order homogeneous, linear, constant coefficients
- Second order nonhomogeneous, constant coefficients, method of undetermined coefficients for certain cases
- Reduction of order, i.e. for $y'' + p(x)y' + q(x)y = 0$, if we have y_1 find y_2 by setting $y_2 = u(x)y_1$.
- Second order nonhomogeneous, variation of parameters.

5 Forced oscillations - resonance, beats, practical resonance

By the end of this section, you should be able to answer the following questions:

- How to determine the steady state solution of a forced oscillator?
- What is resonance?
- How do beats arise?

Recall ODE for free oscillations with damping:

$$my'' + cy' + ky = 0.$$

Now if we have an external force $r(t)$ acting on the body, the equation becomes

$$my'' + cy' + ky = r(t).$$

$r(t)$ is called the input or driving force.

Of particular interest are periodic inputs of the form

$$r(t) = F_0 \cos \omega t, \quad F_0 > 0, \ \omega > 0,$$

so that the ODE becomes

$$my'' + cy' + ky = F_0 \cos \omega t. \tag{6}$$

We have already seen how to determine y_H.

To determine y_P, by the method of undetermined coefficients, we set

$$y_P = a \cos \omega t + b \sin \omega t.$$

After substituting into (6), also setting $\omega_0 = \sqrt{k/m}$, we obtain

$$a = F_0 \frac{m(\omega_0^2 - \omega^2)}{m^2(\omega_0^2 - \omega^2)^2 + \omega^2 c^2} \tag{7}$$

$$b = F_0 \frac{\omega c}{m^2(\omega_0^2 - \omega^2)^2 + \omega^2 c^2} \tag{8}$$

Note we need to modify our initial guess if $\omega = \omega_0$.

We now look at the different cases when the system is damped ($c > 0$) or undamped ($c = 0$).

5.1 Undamped forced oscillations

In this case $c = 0$. Assume $\omega^2 \neq \omega_0^2$. Then

$$y_P = \frac{F_0}{m(\omega_0^2 - \omega^2)} \cos \omega t.$$

We can therefore write the general solution as

$$y(t) = C \cos(\omega_0 t - \delta) + \frac{F_0}{m(\omega_0^2 - \omega^2)} \cos \omega t.$$

This represents a superposition of two harmonic oscillations. Their frequencies are the natural frequency $\omega_0/2\pi$ (cycles/sec) of the system and the frequency $\omega/2\pi$ of the input.

The maximum amplitude of y_P in this case is

$$a_0 = \frac{F_0}{k}\rho, \quad \rho = \frac{1}{1 - (\omega/\omega_0)^2},$$

where ρ is called the resonance factor. As $\omega \to \omega_0$, ρ and $a_0 \to \infty$. This phenomenon of excitation of large oscillations by matching input and natural frequencies ($\omega = \omega_0$) is known as resonance.

In the case of resonance the ODE can be written

$$y'' + \omega_0^2 y = \frac{F_0}{m} \cos \omega_0 t.$$

The modified guess then gives

$$y_P = t(a \cos \omega_0 t + b \sin \omega_0 t).$$

Determining a and b by substitution into the ODE leads to

$$y_P = \frac{F_0}{2m\omega_0} t \sin \omega_0 t.$$

These oscillations grow as t increases.

When we are close to resonance, beats arise.

Take the solution

$$y(t) = \frac{F_0}{m(\omega_0^2 - \omega^2)} (\cos \omega t - \cos \omega_0 t)$$

corresponding to the initial conditions $y(0) = 0$, $y'(0) = 0$. This can be rewritten

$$y(t) = \frac{2F_0}{m(\omega_0^2 - \omega^2)} \sin\left(\frac{\omega_0 + \omega}{2} t\right) \sin\left(\frac{\omega_0 - \omega}{2} t\right)$$

Since we are close to resonance, $\omega_0 - \omega$ is small, so the period of the last sine term is large, giving rise to beats.

5.2 Damped forced oscillations

With damping, $c > 0$ and we know already that

$$y_H = e^{-\frac{c}{2m}t}(A\cos(\psi t) + B\sin(\psi t))$$

(remember underdamping gives damped oscillations).

$y_H \to 0$ as $t \to \infty$, so the general solution in the forced case will approach y_P as $t \to \infty$. That is, the general solution $y(t) = y_H + y_P$ is a transient solution and approaches a steady-state solution which is given by y_P.

This is what happens in practice, because no physical system is completely undamped.

With damping, the amplitude is finite as ω becomes close to ω_0, but may have a large maximum at some value of ω. In otherwords, some input may excite large destructive oscillations even with damping.

For the steady state solution, we have already seen that

$$\begin{aligned} y_P &= a\cos\omega t + b\sin\omega t \\ &= C^*\cos(\omega t - \eta) \end{aligned}$$

with a and b given by (7) and (8) respectively.

The amplitude C^* of y_P is given by

$$\begin{aligned} C^* &= \sqrt{a^2 + b^2} \\ &= \frac{F_0}{\sqrt{m^2(\omega_0^2 - \omega^2)^2 + \omega^2 c^2}} \end{aligned}$$

Treating the amplitude as a function of ω, $C^*(\omega)$ will have a maximum when

$$\frac{dC^*}{d\omega} = 0,$$

that is, when $c^2 - 2m^2(\omega_0^2 - \omega^2) = 0$, or when

$$\omega^2 = \omega_0^2 - \frac{c^2}{2m^2}. \tag{9}$$

For sufficiently large damping, $c^2 > 2m^2\omega_0^2$, (9) has no real solutions, and C^* decreases in a monotone way as ω increases.

If $c^2 \leq 2mk$, (9) has one real solution (remember $\omega > 0$)

$$\omega = \omega_{\max} = \sqrt{\omega_0^2 - \frac{c^2}{2m^2}}$$

and

$$C^*_{\max} = C^*(\omega_{\max}) = \frac{2mF_0}{c\sqrt{4m^2\omega_0^2 - c^2}}.$$

This is what we call practical resonance.

The ratio C^*/F_0 is called the amplification, which $\to \infty$ as $c \to 0$ in agreement with the case of resonance.

6 Hyperbolic functions

By the end of this section, you should be able to answer the following questions:

- What is the definition of the sinh and cosh functions?
- What is the definition of the inverse hyperbolic functions?
- What are the derivatives and anti-derivatives of these functions?
- How are hyperbolic functions used in the catenary problem?

6.1 Properties of hyperbolic functions

We define the functions $\cosh(x)$ and $\sinh(x)$ by

$$\begin{aligned} \cosh(x) &= \frac{e^x + e^{-x}}{2}, \\ \sinh(x) &= \frac{e^x - e^{-x}}{2}. \end{aligned}$$

We can check by direct calculation that

$$\cosh^2(x) - \sinh^2(x) = 1.$$

Compare this with the identity

$$\cos^2(\theta) + \sin^2(\theta) = 1 \tag{10}$$

for trig functions. The identity (10) allows us to parametrise a unit circle. By setting $x(t) = \cos(t)$, $y(t) = \sin(t)$, we have

$$\cos^2(t) + \sin^2(t) = x^2 + y^2 = 1,$$

which is the equation of the unit circle.

If we set $x(t) = \cosh(t)$ and $y(t) = \sinh(t)$, this gives a parametrisation for a hyperbola (only the right branch), since

$$\cosh^2(t) - \sinh^2(t) = x^2 - y^2 = 1,$$

which is the equation of a hyperbola. This is why we call these functions "hyperbolic functions".

These hyperbolic functions satisfy properties similar to their trigonometric counterparts. For example

$$\begin{aligned}\frac{d}{dx}(\cosh(x)) &= \frac{e^x - e^{-x}}{2} = \sinh(x),\\ \frac{d}{dx}(\sinh(x)) &= \frac{e^x + e^{-x}}{2} = \cosh(x).\end{aligned}$$

$\cosh(0) = 1$, $\cosh(x) \geq 1$, $\cosh(x)$ is an even function.

$\sinh(0) = 0$, $\sinh(x)$ is an odd function.

We also define

$$\begin{aligned}\tanh(x) &= \frac{\sinh(x)}{\cosh(x)} = \frac{1 - e^{-2x}}{1 + e^{-2x}}, \quad |\tanh(x)| < 1,\\ \coth(x) &= \frac{\cosh(x)}{\sinh(x)}.\end{aligned}$$

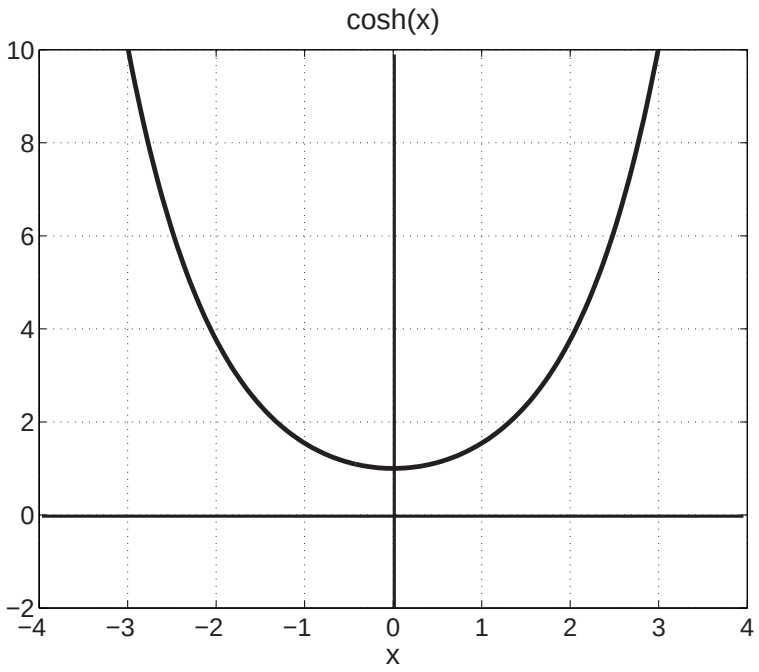

Figure 1: Graph of $\cosh(x)$

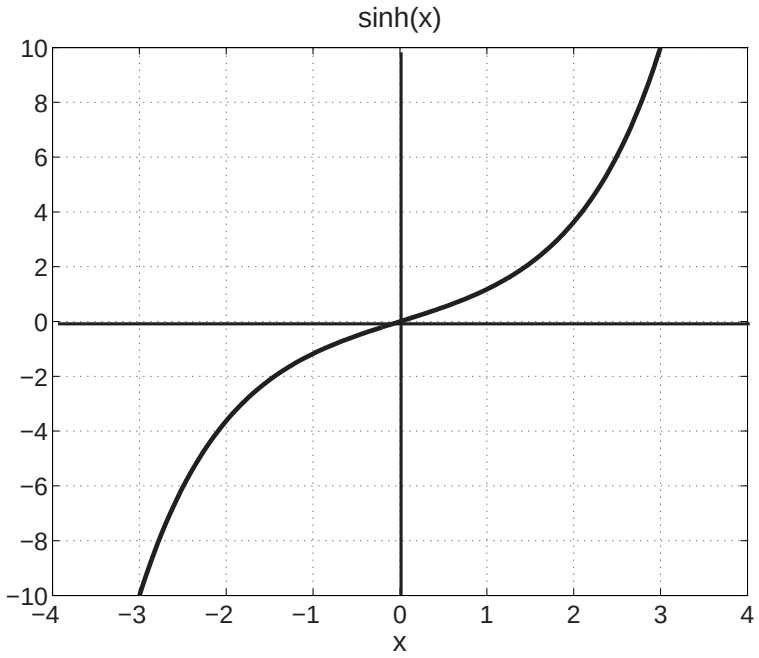

Figure 2: Graph of $\sinh(x)$

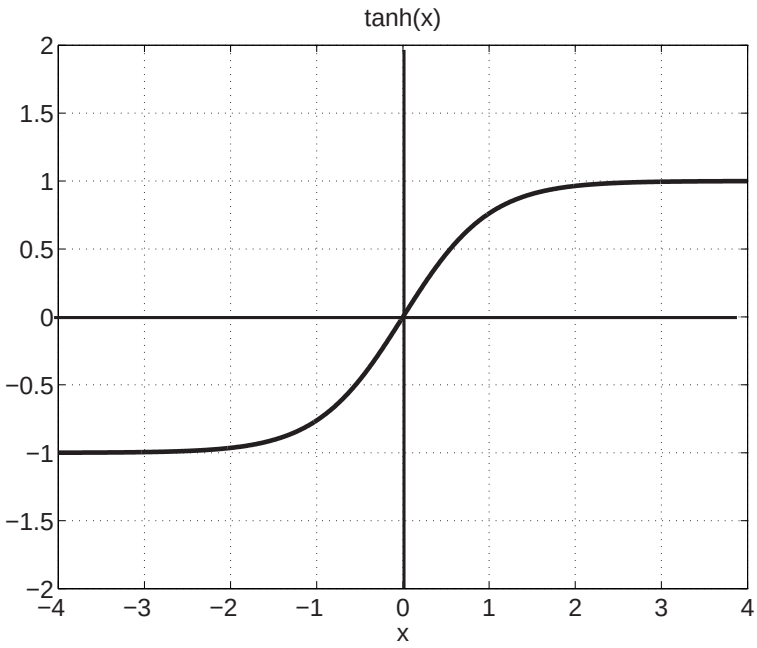

Figure 3: Graph of $\tanh(x)$

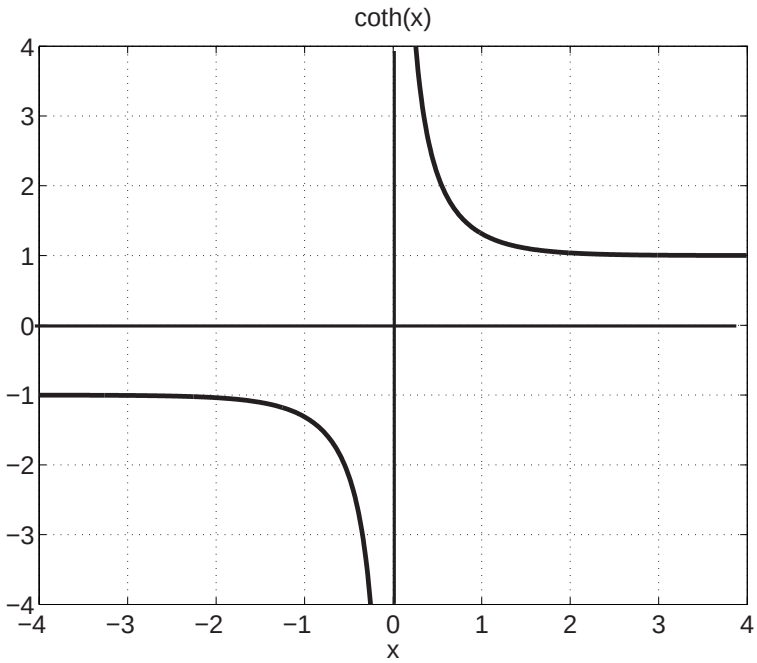

Figure 4: Graph of $\coth(x)$

6.2 Inverse hyperbolic functions

The inverse function of cosh is denoted arcosh .

The inverse function of sinh is denoted arsinh .

The inverse function of tanh is denoted artanh .

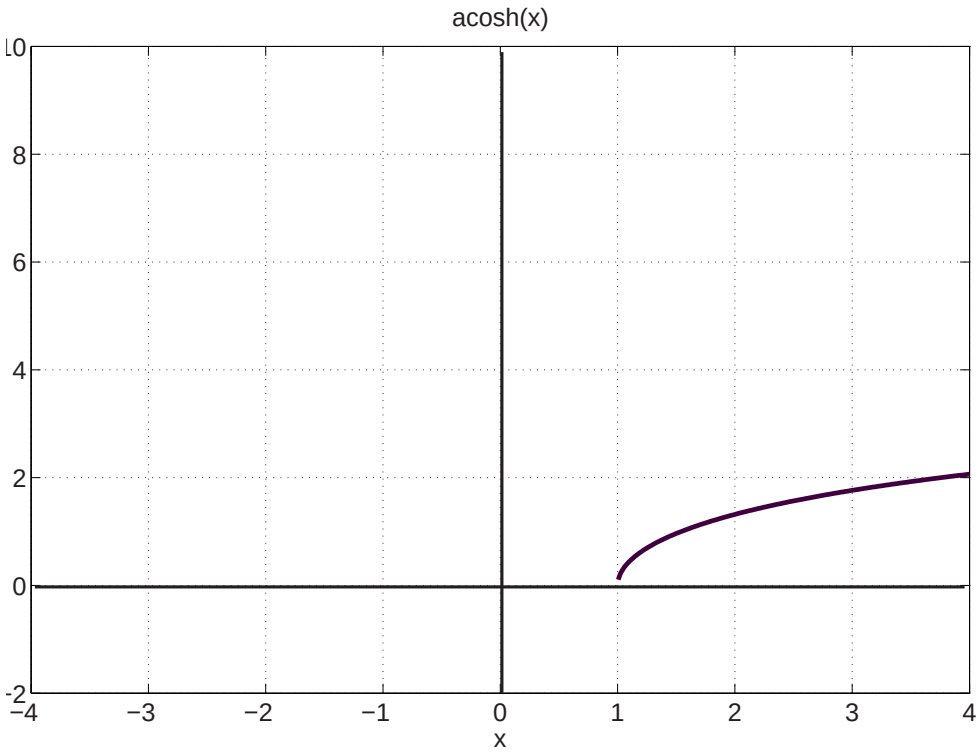

Figure 5: Graph of arcosh (x)

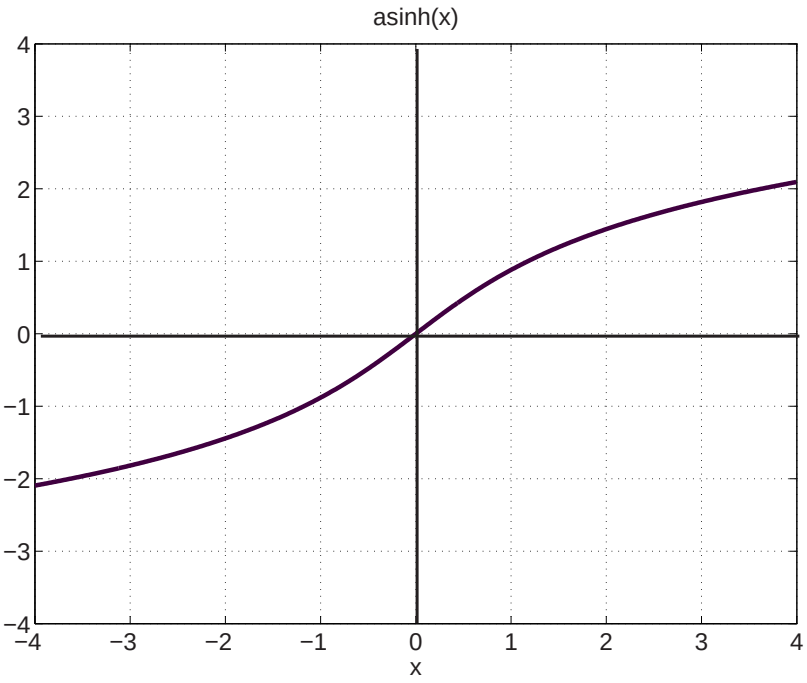

Figure 6: Graph of $\operatorname{arsinh}(x)$

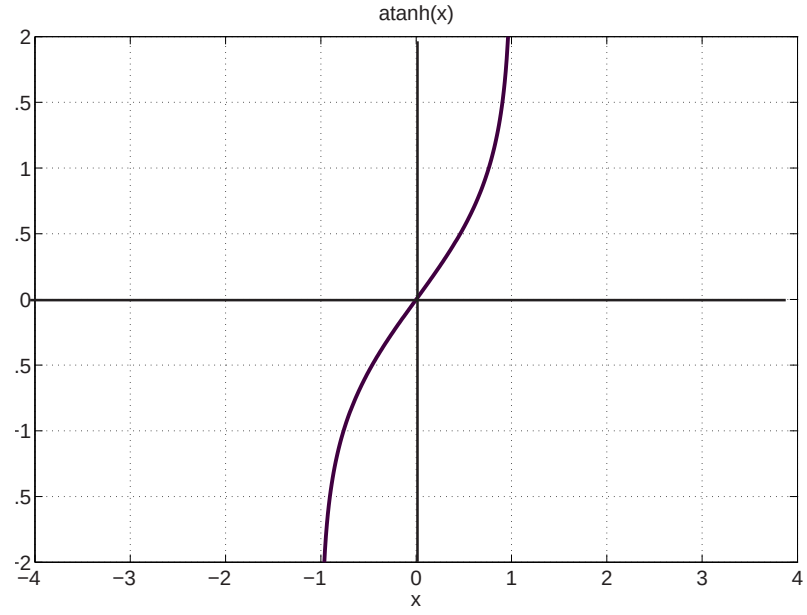

Figure 7: Graph of $\operatorname{artanh}(x)$

We have the following:

$$\begin{aligned}\int \frac{dx}{\sqrt{1+x^2}} &= \operatorname{arsinh}(x) + c \\ \int \frac{dx}{\sqrt{x^2-1}} &= \operatorname{arcosh}(x) + c,\ x > 1.\end{aligned}$$

6.2.1 Show that $\frac{d}{dx}(\mathbf{arsinh}(x)) = \frac{1}{\sqrt{1+x^2}}$

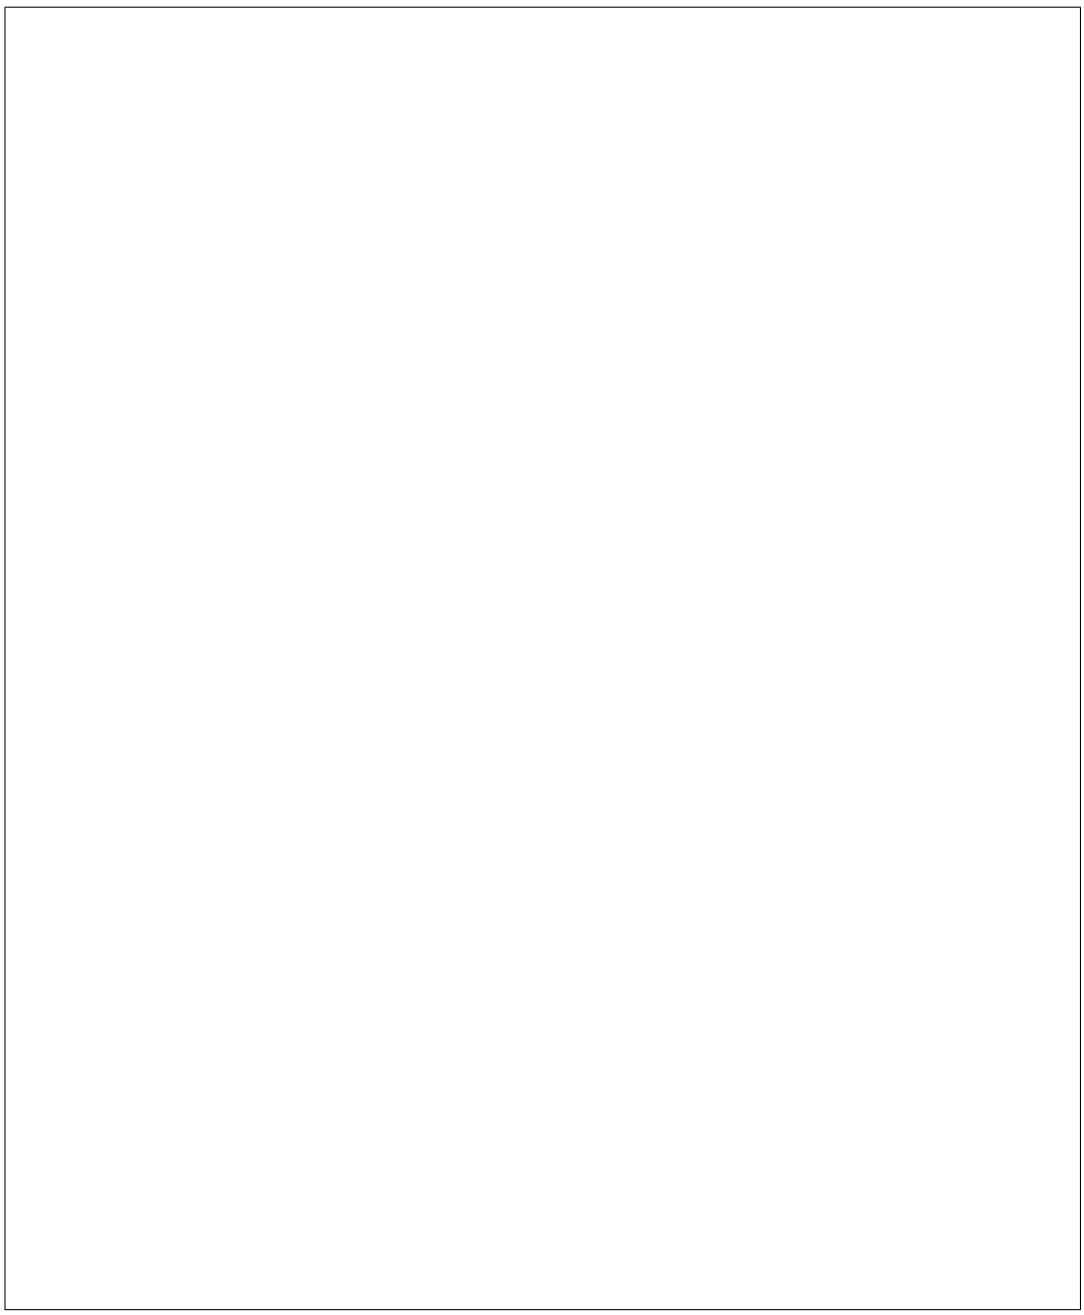

6.2.2 Evaluate the integrals $\int \frac{dx}{\sqrt{1+x^2}}$ and $\int \frac{dx}{\sqrt{x^2-1}}$

6.2.3 Show that $\frac{d}{dx}\left(\mathbf{artanh}\,(x)\right) = \frac{1}{1-x^2}, \quad |x| < 1$

Using partial fractions, we can also evaluate the integral

$$\int \frac{dx}{1-x^2} = \frac{1}{2}\ln\left(\frac{1+x}{1-x}\right) + C.$$

In fact, we have the following identities

$$\begin{aligned}
\operatorname{artanh}(x) &= \frac{1}{2}\ln\left(\frac{1+x}{1-x}\right), \\
\operatorname{arsinh}(x) &= \ln\left(x+\sqrt{x^2+1}\right), \\
\operatorname{arcosh}(x) &= \ln\left(x+\sqrt{x^2-1}\right).
\end{aligned}$$

6.2.4 Show that $\textbf{arsinh}\,(x) = \ln\left(x + \sqrt{x^2+1}\right)$

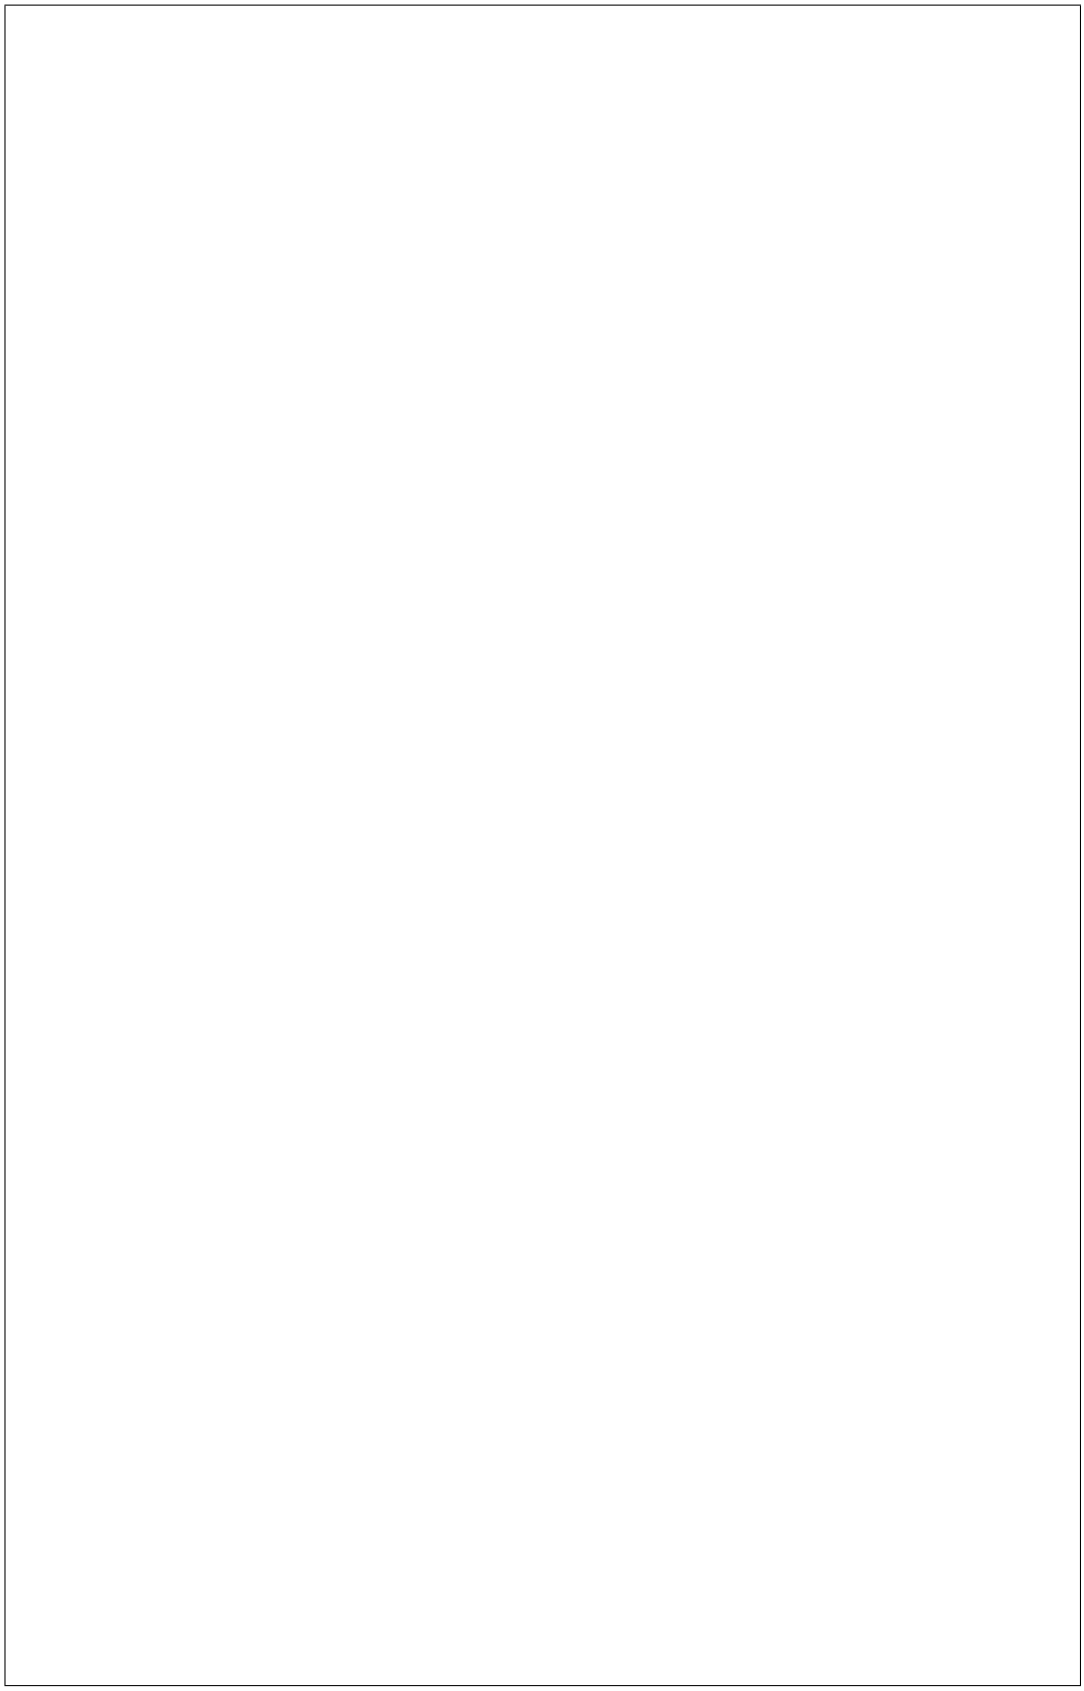

6.3 Reading: The catenary problem

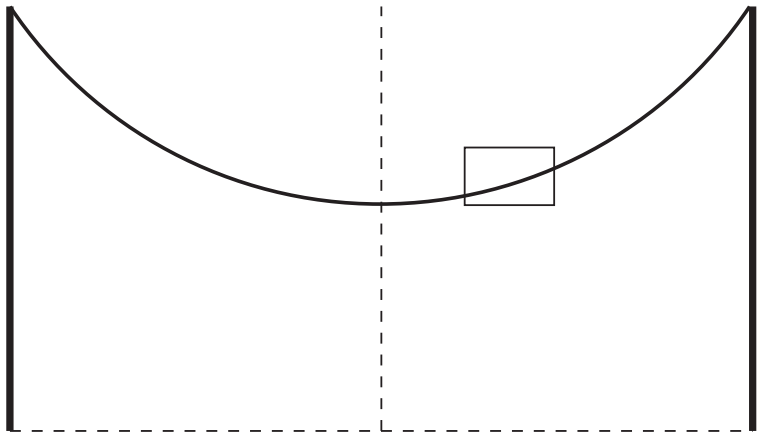

Figure 8: Profile of a heavy chain hanging under gravity.

One of the most famous problems where hyperbolic functions are used is in determining the profile of a heavy chain (of constant density ρ) suspended from two points of equal height (known as a catenary curve).

To derive the differential equation satisfied by the profile $y(x)$, we look at the forces acting on a small element of arc (inside the rectangular box in figure 8).

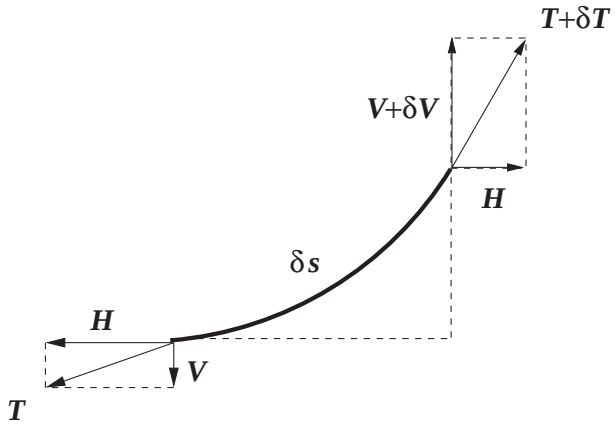

Figure 9: A small arc of heavy chain of length δs.

Let $T(x)$ be the tensile force in the chain with constant horizontal component H (since the load has no x component) and vertical component $V(x)$. In figure 9 the vertical components of the tensile force at either end of the arc are V and $V + \delta V$.

The mass of the arc will be $\rho(\delta s)$, so that the force due to gravity is $\rho g(\delta s)$.

The horizontal equilibrium is the trivial relation $H = H$, whereas the vertical equilibrium is the more informative

$$(V + \delta V) = V + \rho g(\delta s).$$

Dividing both sides by δx gives

$$\frac{\delta V}{\delta x} = \rho g \frac{\delta s}{\delta x}.$$

From geometry, we also have the approximation

$$\frac{\delta y}{\delta x} \approx \frac{V}{H}.$$

We also have the approximation to the arclength δs

$$(\delta s)^2 \approx (\delta x)^2 + (\delta y)^2 \quad \Rightarrow \quad \frac{\delta s}{\delta x} \approx \sqrt{1 + \left(\frac{\delta y}{\delta x}\right)^2}$$

Finally we take the limit $\delta x \to 0$ so that $\delta y \to 0$ and $\delta s \to 0$ simultaneously. We then have the following equations

$$\begin{aligned} \frac{dV}{dx} &= \rho g \frac{ds}{dx}, \\ V &= H \frac{dy}{dx}, \\ \frac{ds}{dx} &= \sqrt{1 + \left(\frac{dy}{dx}\right)^2}. \end{aligned}$$

Putting these equations together gives the ODE satisfied by the profile $y(x)$,

$$\frac{d^2y}{dx^2} = \frac{\rho g}{H} \sqrt{1 + \left(\frac{dy}{dx}\right)^2}.$$

7 Introduction to double integrals, volume below a surface

By the end of this section, you should be able to answer the following questions:

- What is the definition of volume below a surface?
- What is the definition of a double integral?
- How are the two related?

Recall that if $y = f(x)$, the area under the curve over the interval $I = [a, b]$ is

$$\int_I f(x)dx = \lim \sum_{i=1}^{n} f(x_i^*)(x_i - x_{i-1})$$

where $x_i^* \in [x_i, x_{i-1}]$.

7.1 Double integrals

Suppose we have a surface $z = f(x, y)$ above a planar region R in the x-y plane.

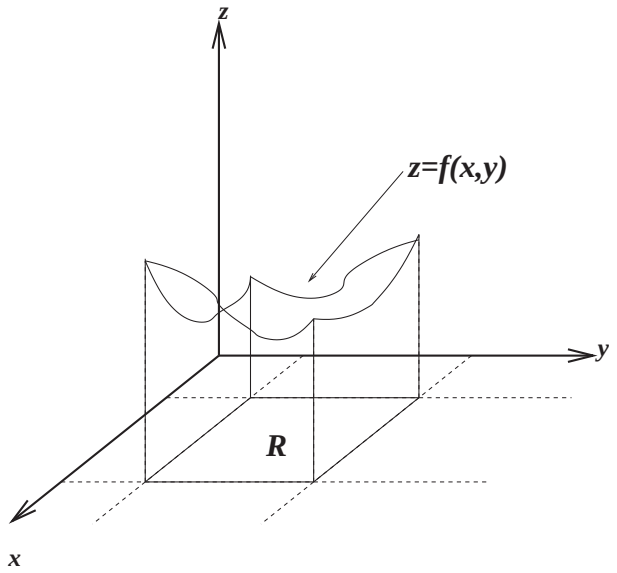

Figure 10: What is the volume V under the surface?

Before moving onto general regions, we start by considering the case where R is a rectangle. That is,

$$R = \{(x, y) \in \mathbb{R}^2 \mid a \leq x \leq b,\ c \leq y \leq d\}.$$

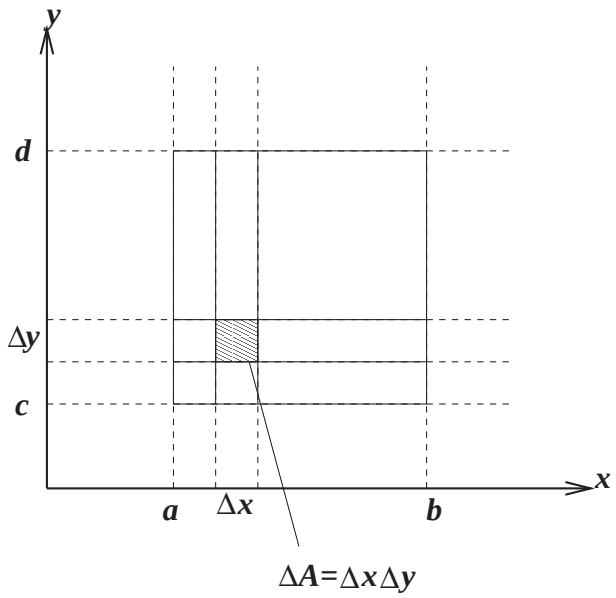

Start by dividing R into subrectangles by dividing the interval $[a, b]$ into m subintervals $[x_{i-1}, x_i]$, each of width $\Delta x = \dfrac{b-a}{m}$ and $[c, d]$ into n subintervals $[y_{j-1}, y_j]$ of equal width $\Delta y = \dfrac{d-c}{n}$.

Combining these gives a rectangular grid R_{ij} with subrectangles each of area $\Delta A = \Delta x \Delta y$.

In each subrectangle take any point P_{ij} with co-ordinates (x^*_{ij}, y^*_{ij}).

The volume of the box with base the rectangle ΔA and height the value of the function $f(x, y)$ at the point P_{ij} (so the box touches the surface at a point directly above P_{ij} - see figure 11) is

$$V_{ij} = f(x^*_{ij}, y^*_{ij})\Delta A.$$

Then for all the subrectangles we have an approximation to the required volume V:

$$V \approx \sum_{i=1}^{m} \sum_{j=1}^{n} f(x^*_{ij}, y^*_{ij})\Delta A,$$

the double Riemann sum.

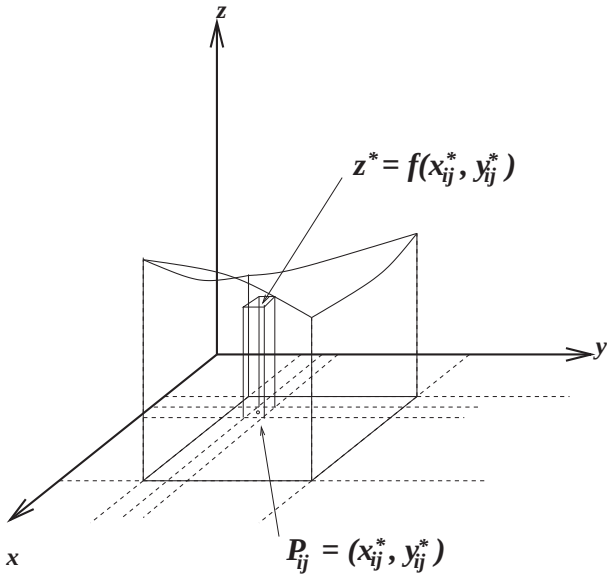

Figure 11: The rectangular box whose volume is $z^*\Delta A$.

Let $\Delta x \to 0$ and $\Delta y \to 0$, ie $m \to \infty$ and $n \to \infty$, then we *define* the volume to be

$$V = \lim_{m\to\infty} \lim_{n\to\infty} \sum_{i=1}^{m} \sum_{j=1}^{n} f(x_{ij}^*, y_{ij}^*)\Delta A,$$

if the limits exist and we write this as

$$\iint_R f(x,y)dA.$$

We call f integrable if the limits exist. Note that every continuous function is integrable.

7.2 Properties of the double integral

(i) $\iint_R (f \pm g)dA = \iint_R fdA \pm \iint_R gdA$

(ii) $\iint_R cfdA = c\iint_R fdA$

(iii) $\iint_R fdA = \iint_{R_1} fdA + \iint_{R_2} fdA$

(iv) If $f(x,y) \geq g(x,y)$ for all $(x,y) \in R$ then

$$\iint_R fdA \geq \iint_R gdA$$

7.3 Iterated integrals

We define $\int_c^d f(x,y)dy$ to mean that x is fixed and $f(x,y)$ is integrated with respect to y from $y = c$ to $y = d$. So

$$A(x) = \int_c^d f(x,y)dy$$

is a function of x only.

If we now integrate $A(x)$ with respect to x from $x = a$ to $x = b$ we have

$$\begin{aligned} \int_a^b A(x)dx &= \int_a^b \left[\int_c^d f(x,y)dy\right] dx \\ &= \int_a^b \int_c^d f(x,y)\ dy\ dx \end{aligned}$$

This is called an iterated integral.

7.3.1 Example: evaluate $\int_0^2 \int_1^3 x^2 y\ dy\ dx$

Now try integrating the other way around:

7.3.2 Example: evaluate $\int_1^3 \int_0^2 x^2 y \; dx \; dy$

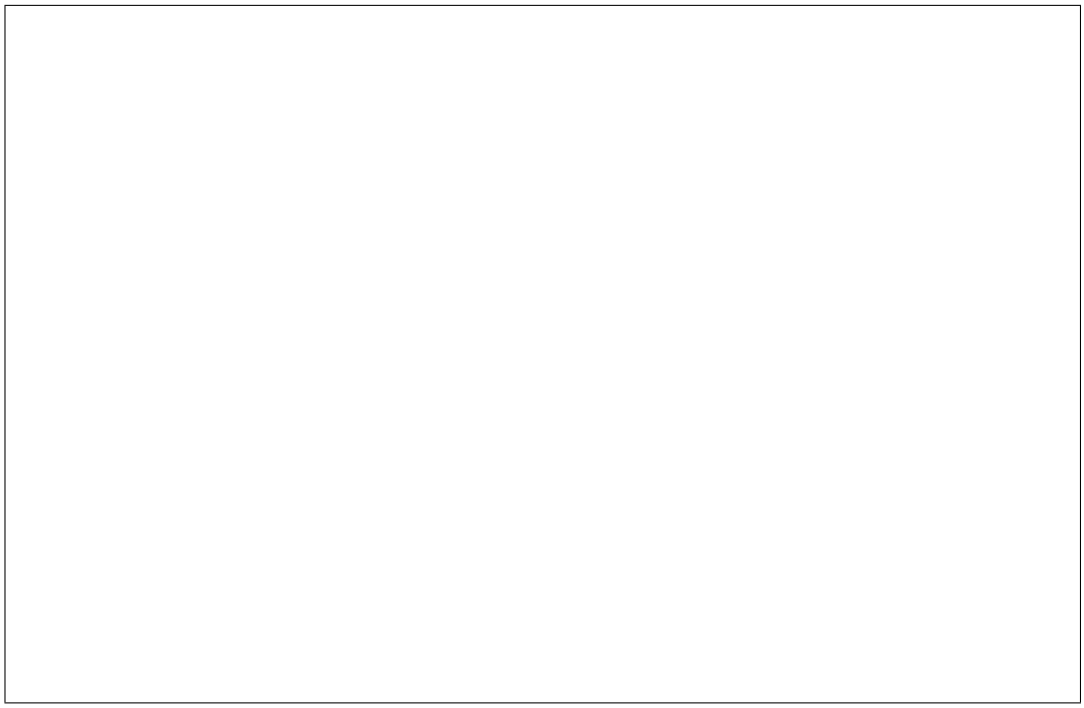

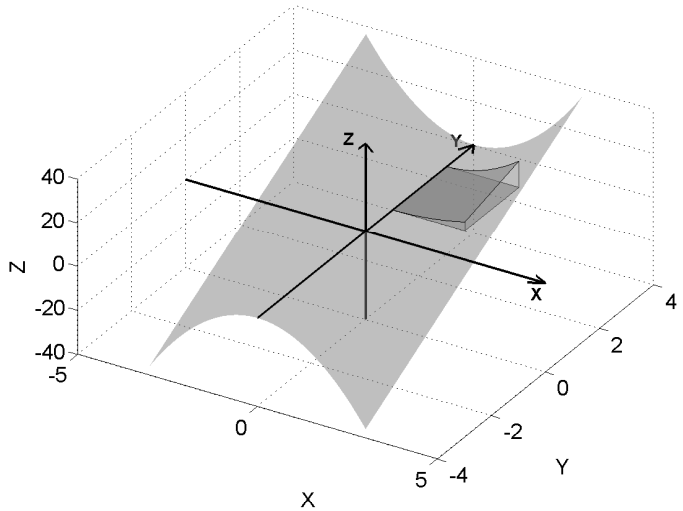

Figure 12: We have just calculated the volume of the solid outlined above.

8 Fubini's theorem, volume by slabs

By the end of this section, you should be able to answer the following questions:

- What is Fubini's theorem?
- How is the double integral related to the iterated integral?
- How do you estimate the volume below a surface using slabs?

8.1 Fubini's theorem

If $f(x, y)$ is integrable on the rectangle

$$R = \{(x, y) | a \leq x \leq b,\ c \leq y \leq d\},$$

then

$$\begin{aligned} \iint_R f(x, y) dA &= \int_a^b \int_c^d f(x, y)\ dy\ dx \\ &= \int_c^d \int_a^b f(x, y)\ dx\ dy \end{aligned}$$

8.2 Example: evaluate $\iint_R (x^2 + y^2) dA$ where $R = \{(x, y) | 0 \leq x \leq 2,\ 0 \leq y \leq 1\}$

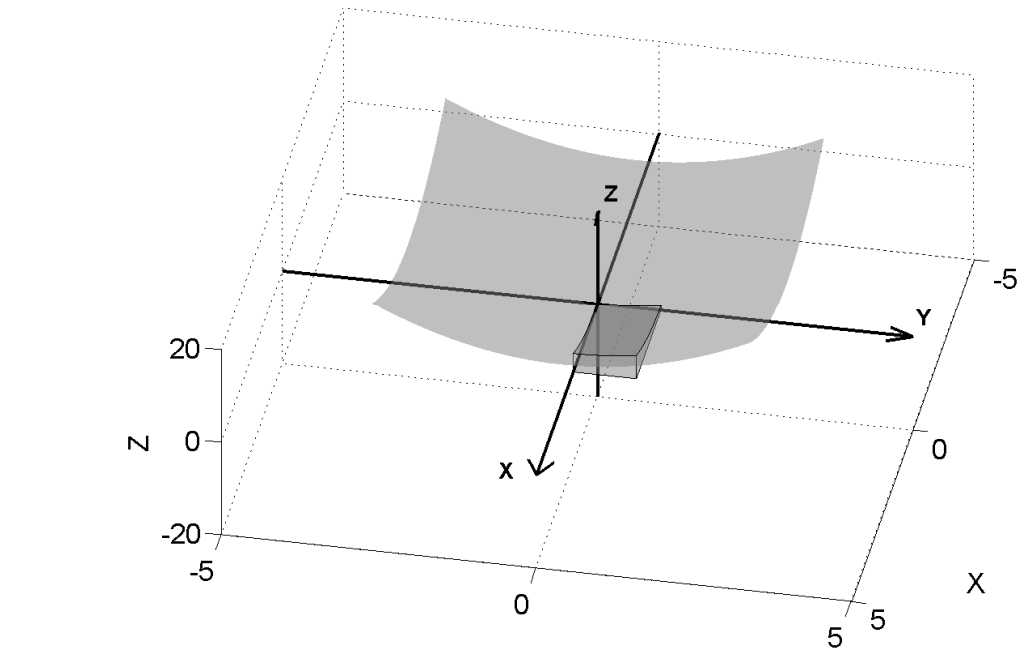

Figure 13: A representation of the volume in example 8.2.

8.3 Interpreting Fubini's theorem in terms of volume

Fubini's theorem is the key result that tells us how to evaluate a double integral. We can see the relation between the iterated integral and the double integral by considering an alternative way of calculating the volume below a surface.

Suppose we want to find the volume below the surface $z = x^2y$ above the square region $0 \leq x \leq 8$ and $0 \leq y \leq 4$.

A natural way to solve this problem is to break the region up into slabs of equal depth $\Delta y = y_{j+1} - y_j$ located at y_j, and add up the volume of the slabs

$$V \approx \sum_j \Delta V,$$

where ΔV the volume of the jth slab. Figure 14 below shows two ways of doing this using four slabs in each case. The left diagram follows the method outlined here, taking slabs of thickness Δy.

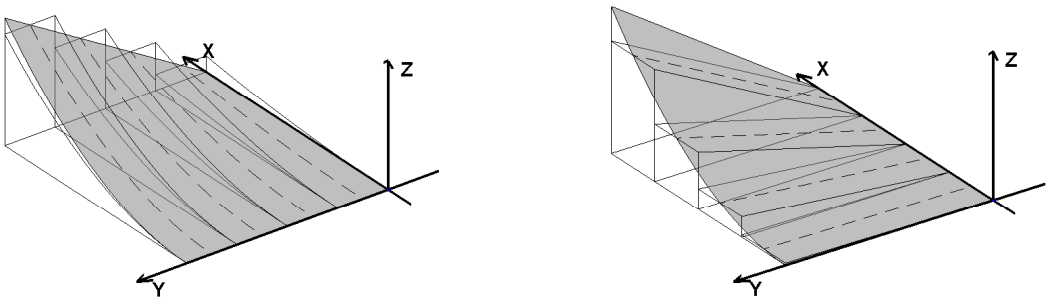

Figure 14: Two ways of approximating the volume under $z = x^2y$ using four slabs.

If the slab is very thin (i.e. $\Delta y \ll 1$) then the volume of each slab is

$$\Delta V \approx \text{ Area of slab } \times \text{ Depth } = C(y_j)\Delta y.$$

Here $C(y_j)$ is the area of the slab at the location y_j (and the result will depend on y_j!). From one-dimensional calculus we know exactly that

$$C(y_j) = \int_0^8 f(x, y_j)dx \quad y_j \text{ constant.}$$

It is easy to compute this as a regular integral since y_j does not vary with x. Putting all this together

$$V \approx \sum_j \Delta V_j \approx \sum_j C(y_j)\Delta y.$$

As the slabs become thinner and thinner ($\Delta y \to 0$) the approximation becomes more accurate and we can replace the summation by an integral[1]

$$V = \int_0^4 C(y)dy = \int_0^4 \left(\int_0^8 f(x,y)dx \right) dy$$

Note that the y is held constant in the inner integral.

A similar argument can be applied by considering slabs of depth Δx, located at x_j. In other words, take slabs that are parallel to the y-z plane.

8.4 Example: find the volume of the solid bounded by the elliptic paraboloid $x^2 + 2y^2 + z = 16$, the planes $x = 2$ and $y = 2$, and the three coordinate planes.

[1]Recall that is in fact the definition of an integral

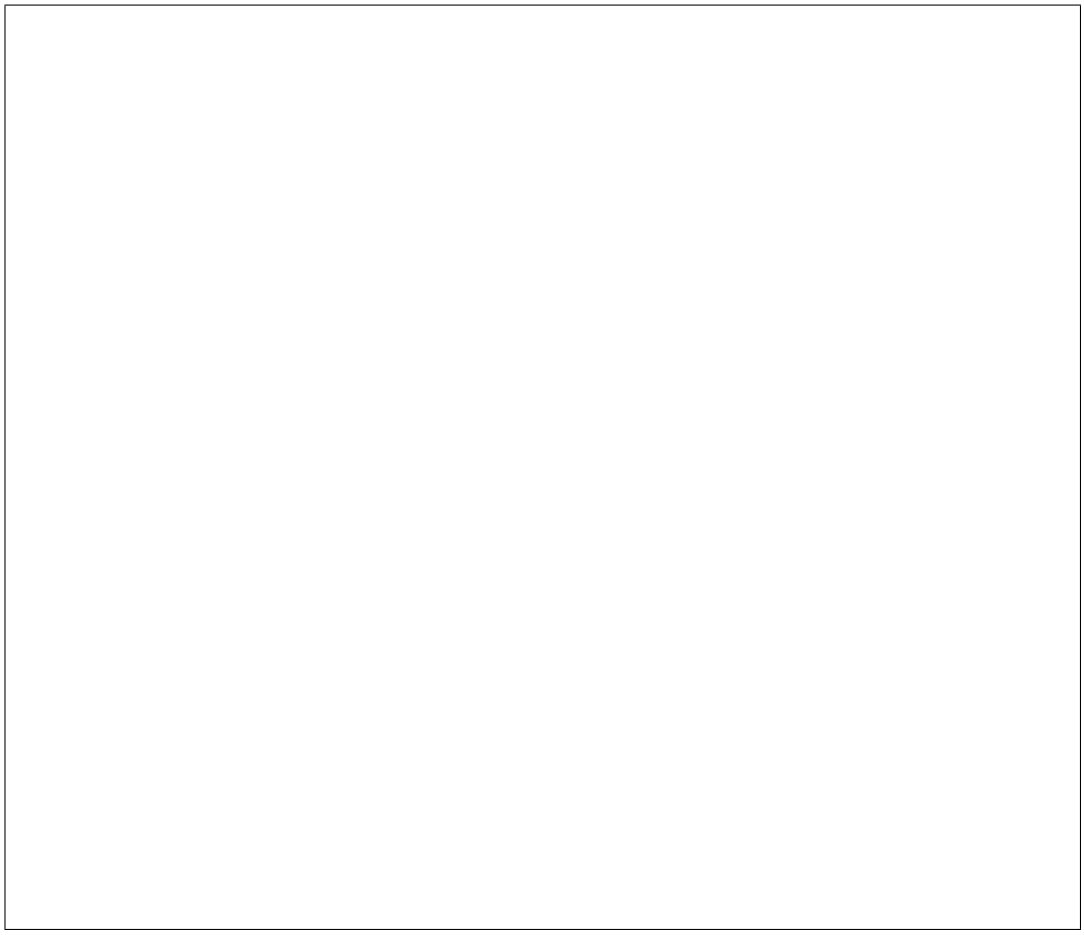

Figure 15: The volume of the solid of example 8.4 is below the surface $z = 16 - x^2 - 2y^2$ and above the x-y plane as shown.

8.5 Special case when $f(x,y) = g(x)h(y)$.

In this case we can separate the integral as follows.

$$\begin{aligned} \iint_R f(x,y)dA &= \int_c^d \int_a^b g(x)h(y)\ dx\ dy \\ &= \int_a^b g(x)dx \int_c^d h(y)dy \end{aligned}$$

8.5.1 Example: $\iint_R \sin x \cos y\ dA$ where $R = [0, \frac{\pi}{2}] \times [0, \frac{\pi}{2}]$

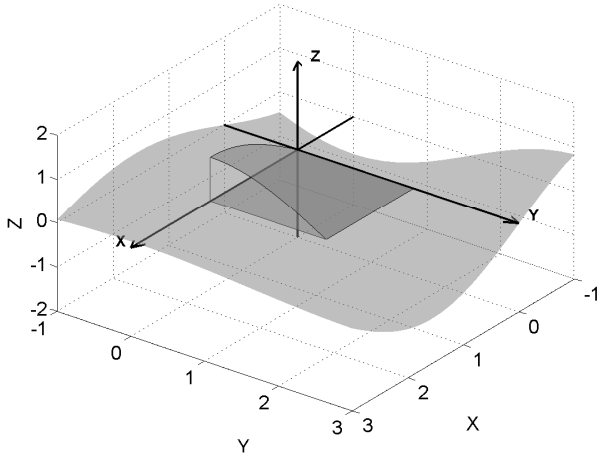

Figure 16: The volume calculated in example 8.5.1 is outlined above.

9 Integrals over general regions

By the end of this section, you should be able to answer the following questions:

- How can you identify type I and II regions?
- How do you evaluate a double integral over type I and II regions?
- How can you evaluate a double integral over a more general region comprising finitely many type I and II regions?
- What is meant by net volume below a surface?

To find the double integral over a general region D instead of just a rectangle we consider a rectangle which encloses D and define

$$F(x,y) = \begin{cases} f(x,y), & \text{if } (x,y) \in D \\ 0, & \text{if } (x,y) \in R \text{ but } \notin D \end{cases}$$

then

$$\iint_D f(x,y)dA = \iint_R F(x,y)dA$$

and we can proceed as before. It is possible to show that F is integrable if the boundary of D is bounded by a finite number of smooth curves of finite length. Note that F may still be discontinuous at the boundary of D.

9.1 Type I regions

A plane region D is of type I if it lies between the graph of two continuous functions of x. That is $D = \{(x, y)|a \leq x \leq b,\ g_1(x) \leq y \leq g_2(x)\}$.

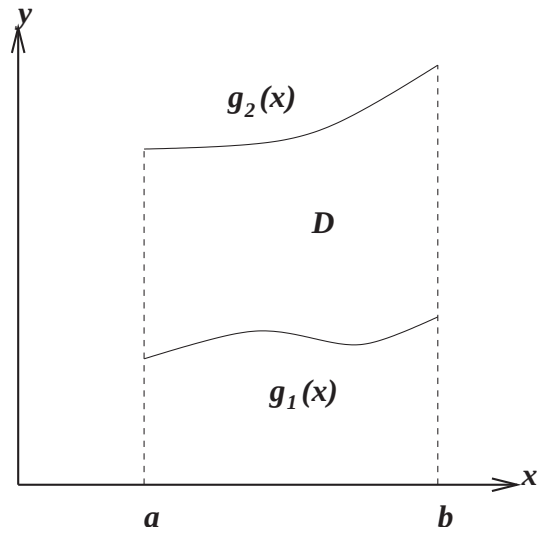

Figure 17: Type I regions are generally bounded by two constant values of x and two functions of x.

In practice, to evaluate $\iint_D f(x, y)dA$ where D is a region of type I we have

$$\iint_D f(x, y)dA = \int_a^b \int_{g_1(x)}^{g_2(x)} f(x, y)\ dy\ dx.$$

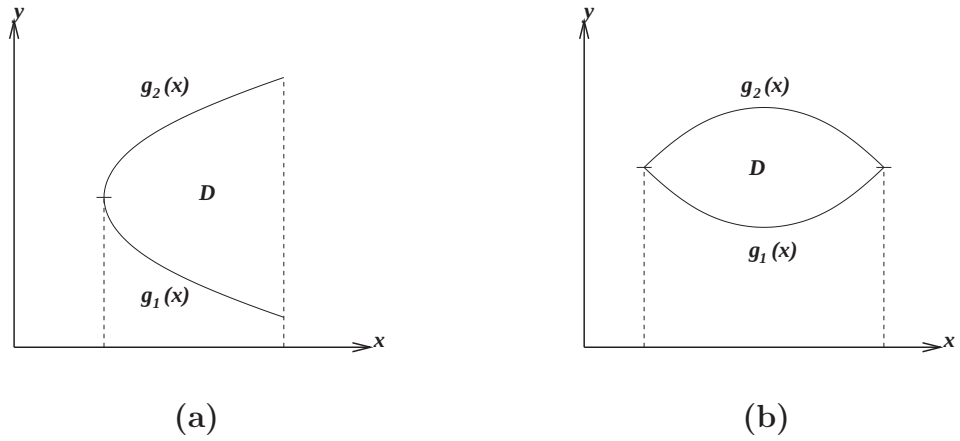

Figure 18: Some more examples of type I regions.

9.1.1 Example: find $\iint_D (4x + 10y)dA$ where D is the region between the parabola $y = x^2$ and the line $y = x + 2$.

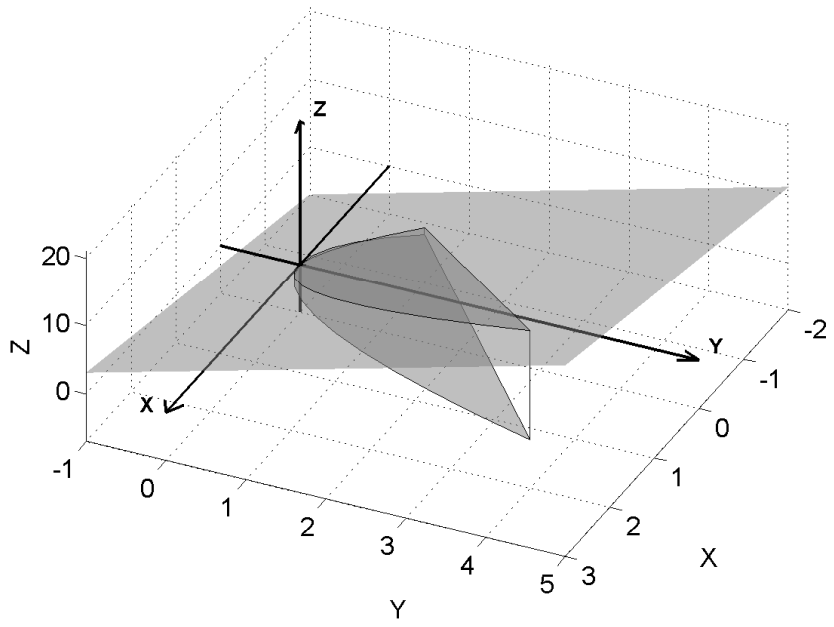

Figure 19: The volume of example 9.1.1 is outlined above.

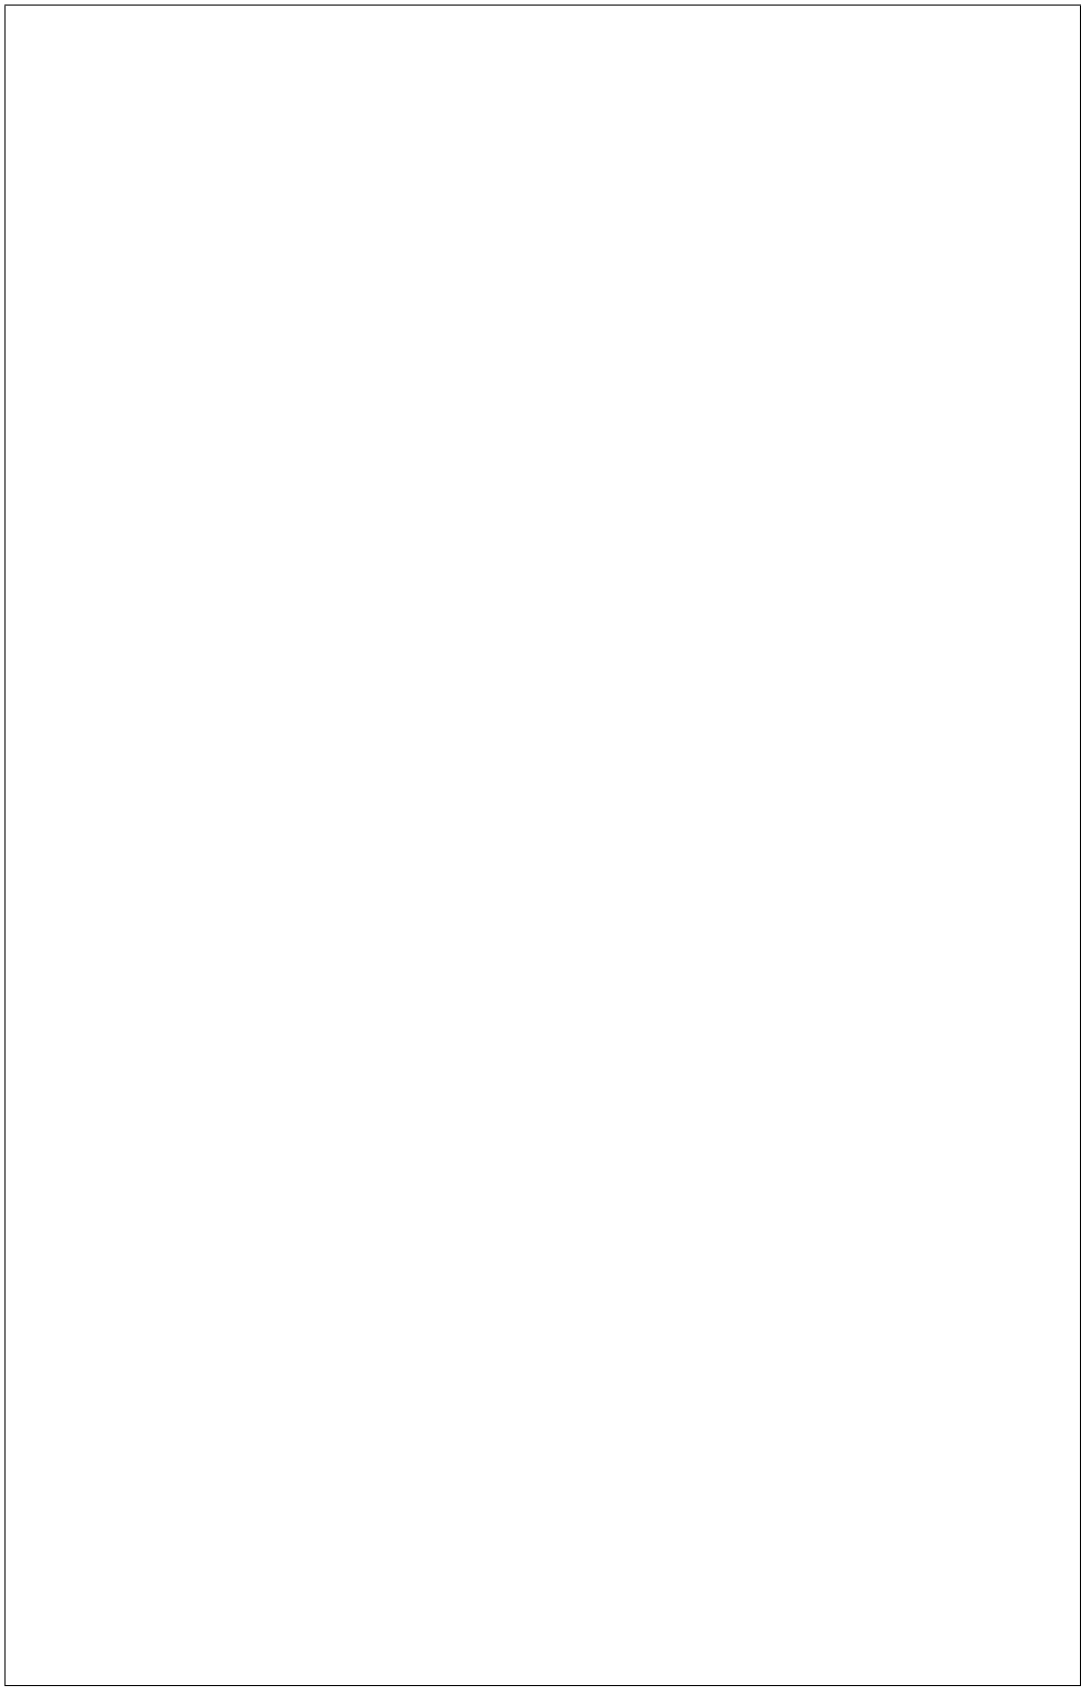

9.2 Type II regions

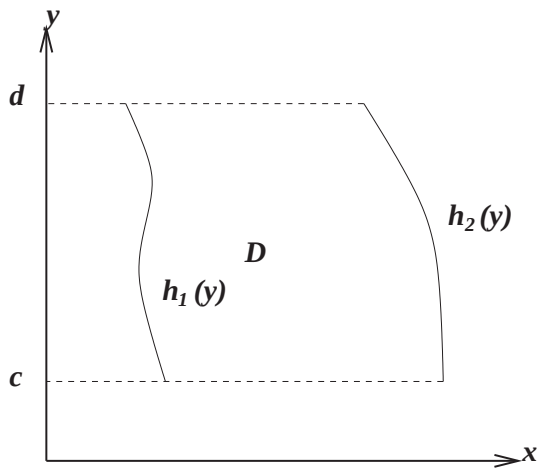

Figure 20: Type II regions are generally bounded by two constant values of y and two functions of y.

A plane region is of type II if it can be expressed by

$$D = \{(x, y) | c \leq y \leq d,\ h_1(y) \leq x \leq h_2(y)\}.$$

In practice, to evaluate $\iint\limits_D f(x, y)dA$ where D is a region of type II we have

$$\iint\limits_D f(x, y)dA = \int_c^d \int_{h_1(y)}^{h_2(y)} f(x, y)\ dx\ dy.$$

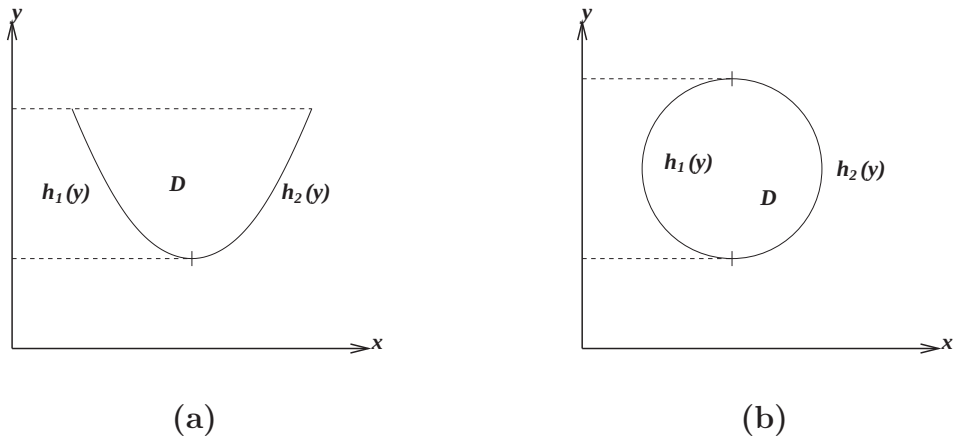

Figure 21: Some more examples of type II regions.

9.2.1 Example: evaluate $\iint_D xy\ dA$ where D is the region bounded by the line $y = x - 1$ and the parabola $y^2 = 2x + 6$.

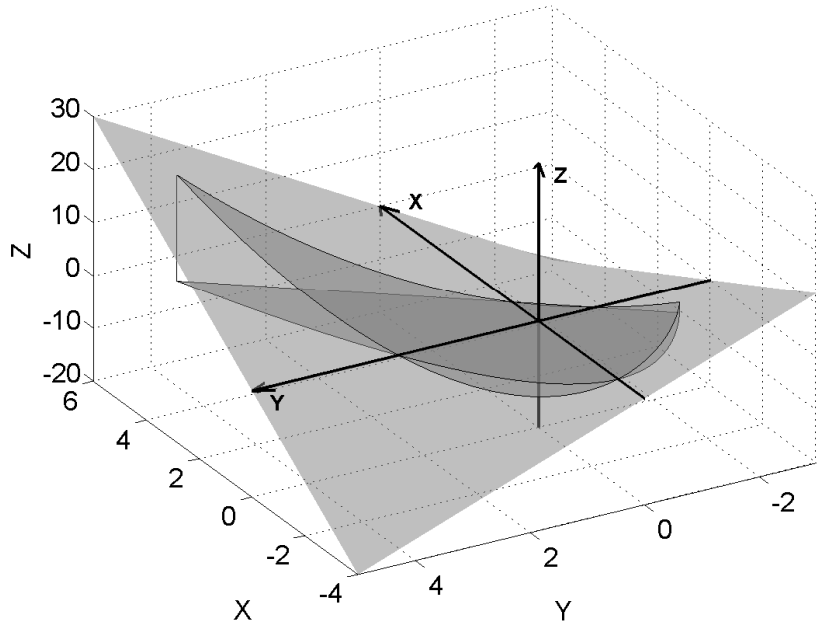

Figure 22: The volume of example 9.2.1 is outlined above. Note carefully that the surface is *above* the x-y plane only in the quadrants where $x, y > 0$ and $x, y < 0$. For x and y values in the other two quadrants, the surface is *below* the x-y plane. Hence in this example we are calculating the "net volume" lying above the x-y plane.

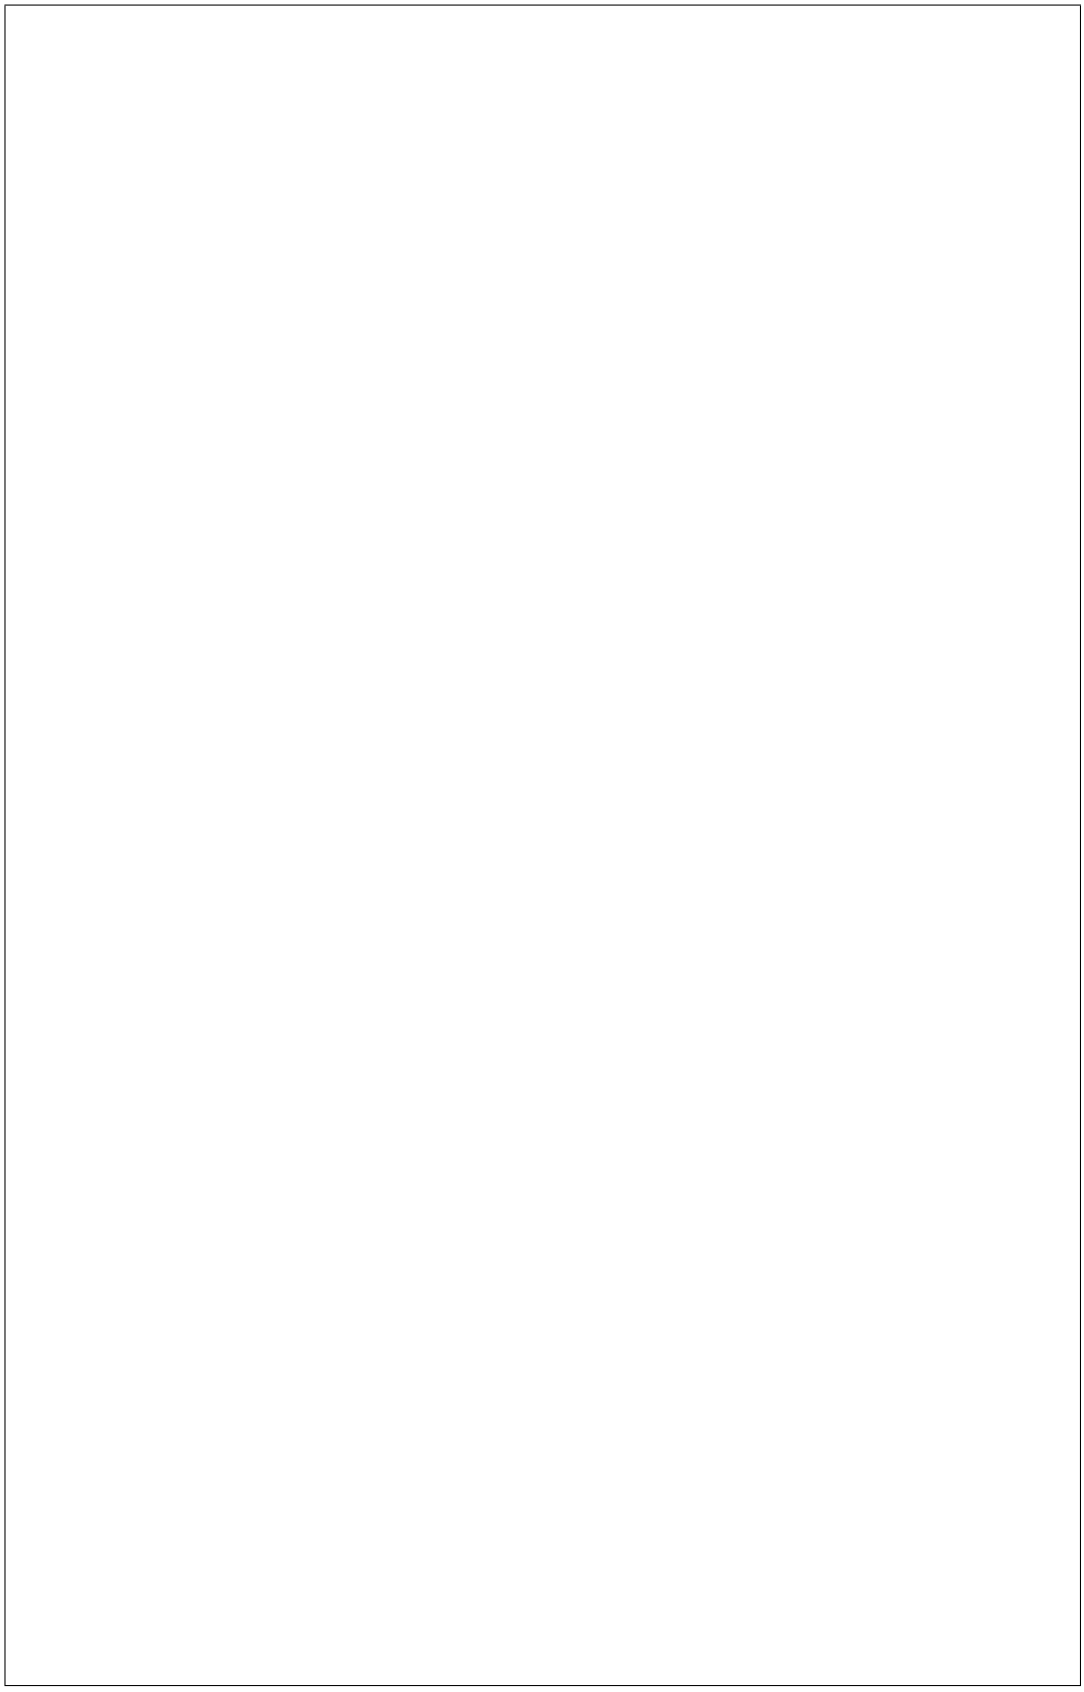

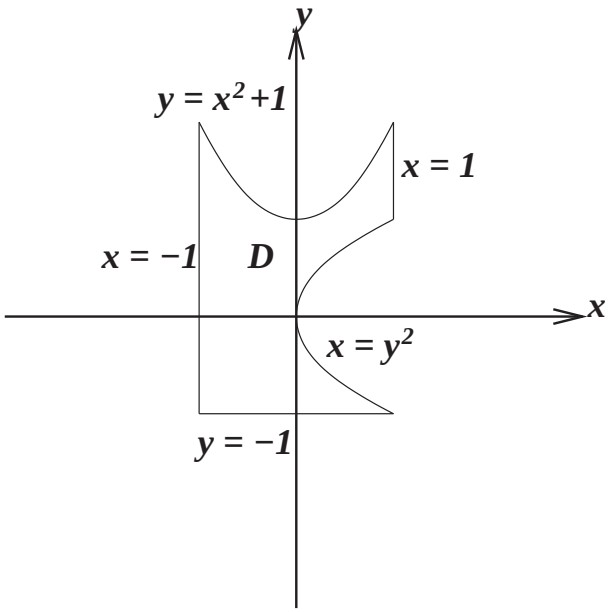

9.3 Express D as a union of regions of type I or type II and expand the integral $\iint_D f(x,y)\ dA$, for some integrable function f.

10 Interchanging order of integration

By the end of this section, you should be able to answer the following questions:

- How do you change the order of integration in a double integral?
- when might it be necessary to change the order of integration in a double integral?

It is often possible to represent a type I region as a union of type II regions, or a type II region as a union of type I regions. Why would we want to do that? In some cases, it may only be possible to integrate a function one way but not the other. In this section, we investigate this idea more closely.

10.1 Find the volume under the paraboloid $z = x^2+y^2$ above the region D, where D is bounded by $y = x^2$ and $y = 2x$. Do the problem twice, first by taking D to be a type I region, then by taking D to be type II.

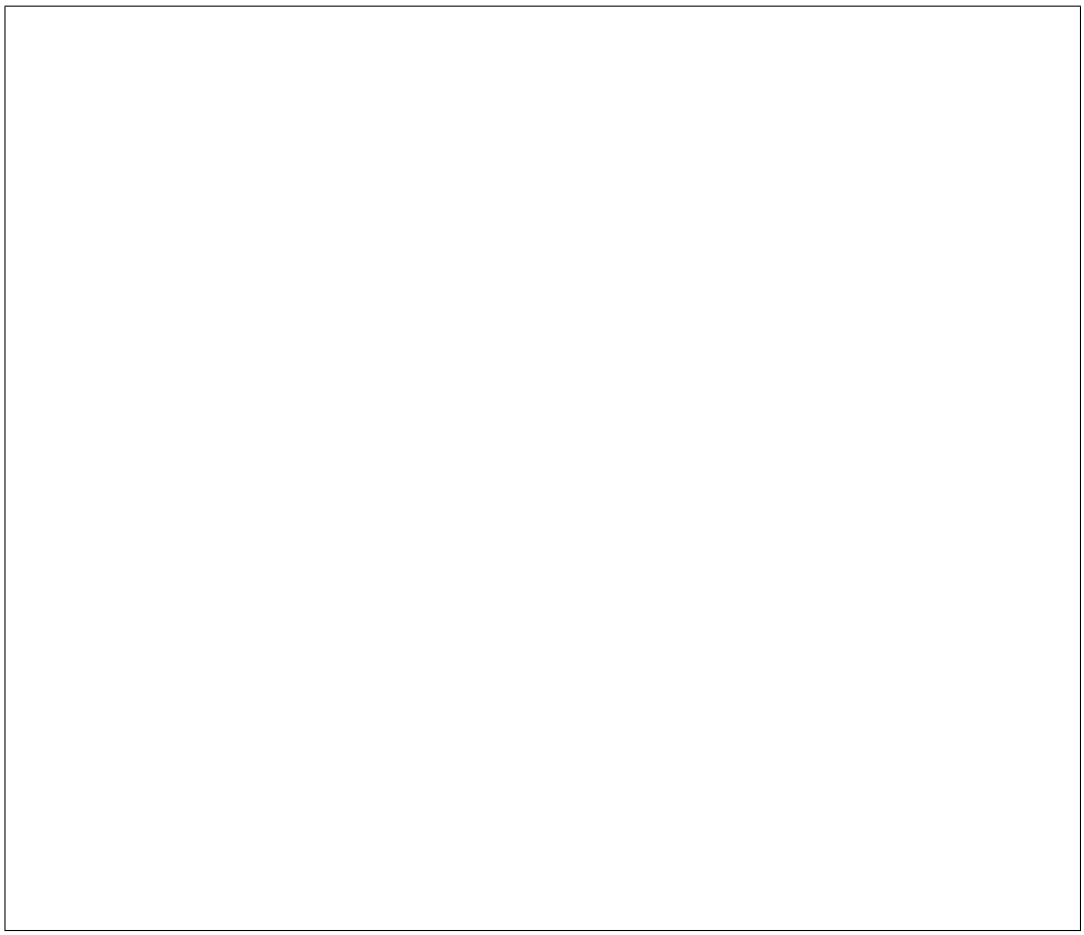

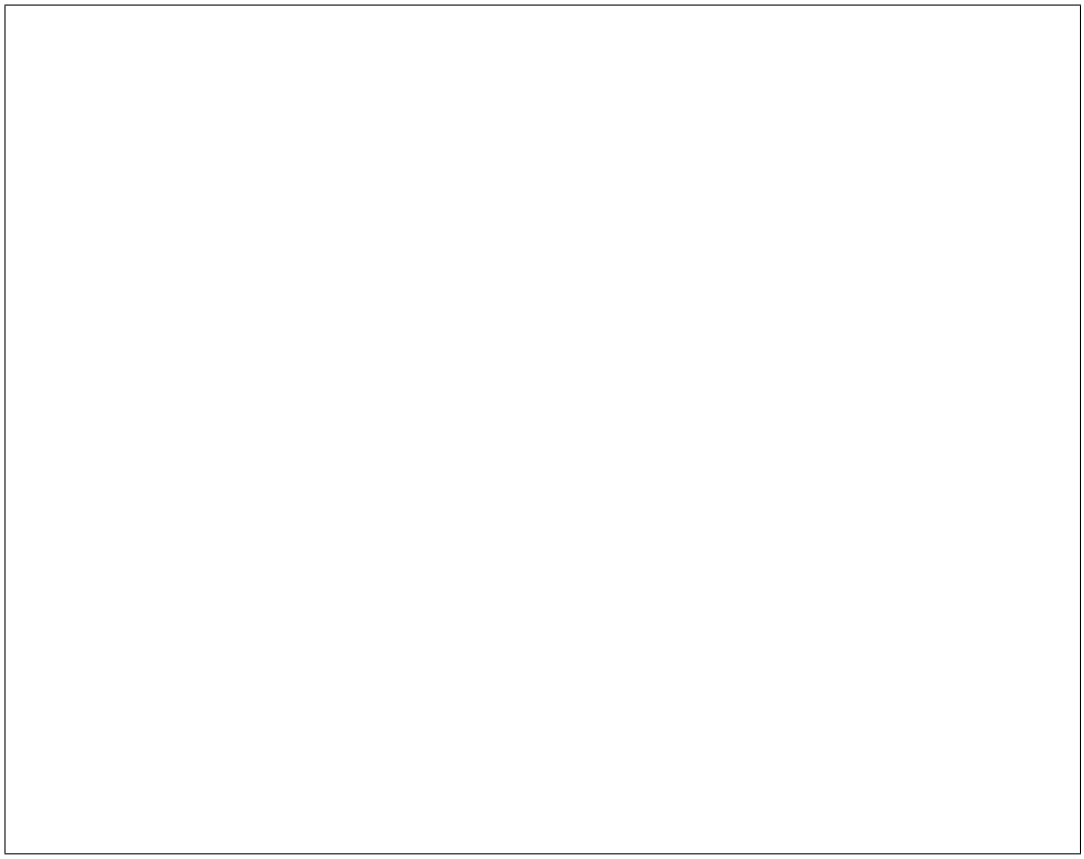

Figure 23: This volume can be calculated by treating the region in the x-y plane as either type I or II as seen in example 10.1.

In the following example, we see how it is sometimes necessary to change the order of integration in order to evaluate the integral.

10.2 Example: Find $\int_0^1 \int_x^1 \sin(y^2)\ dy\ dx$

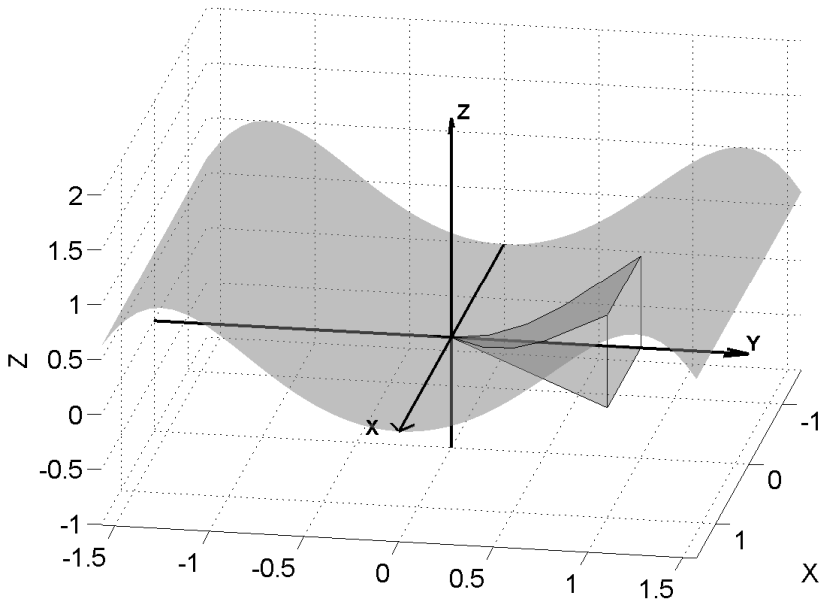

Figure 24: The volume described in example 10.2.

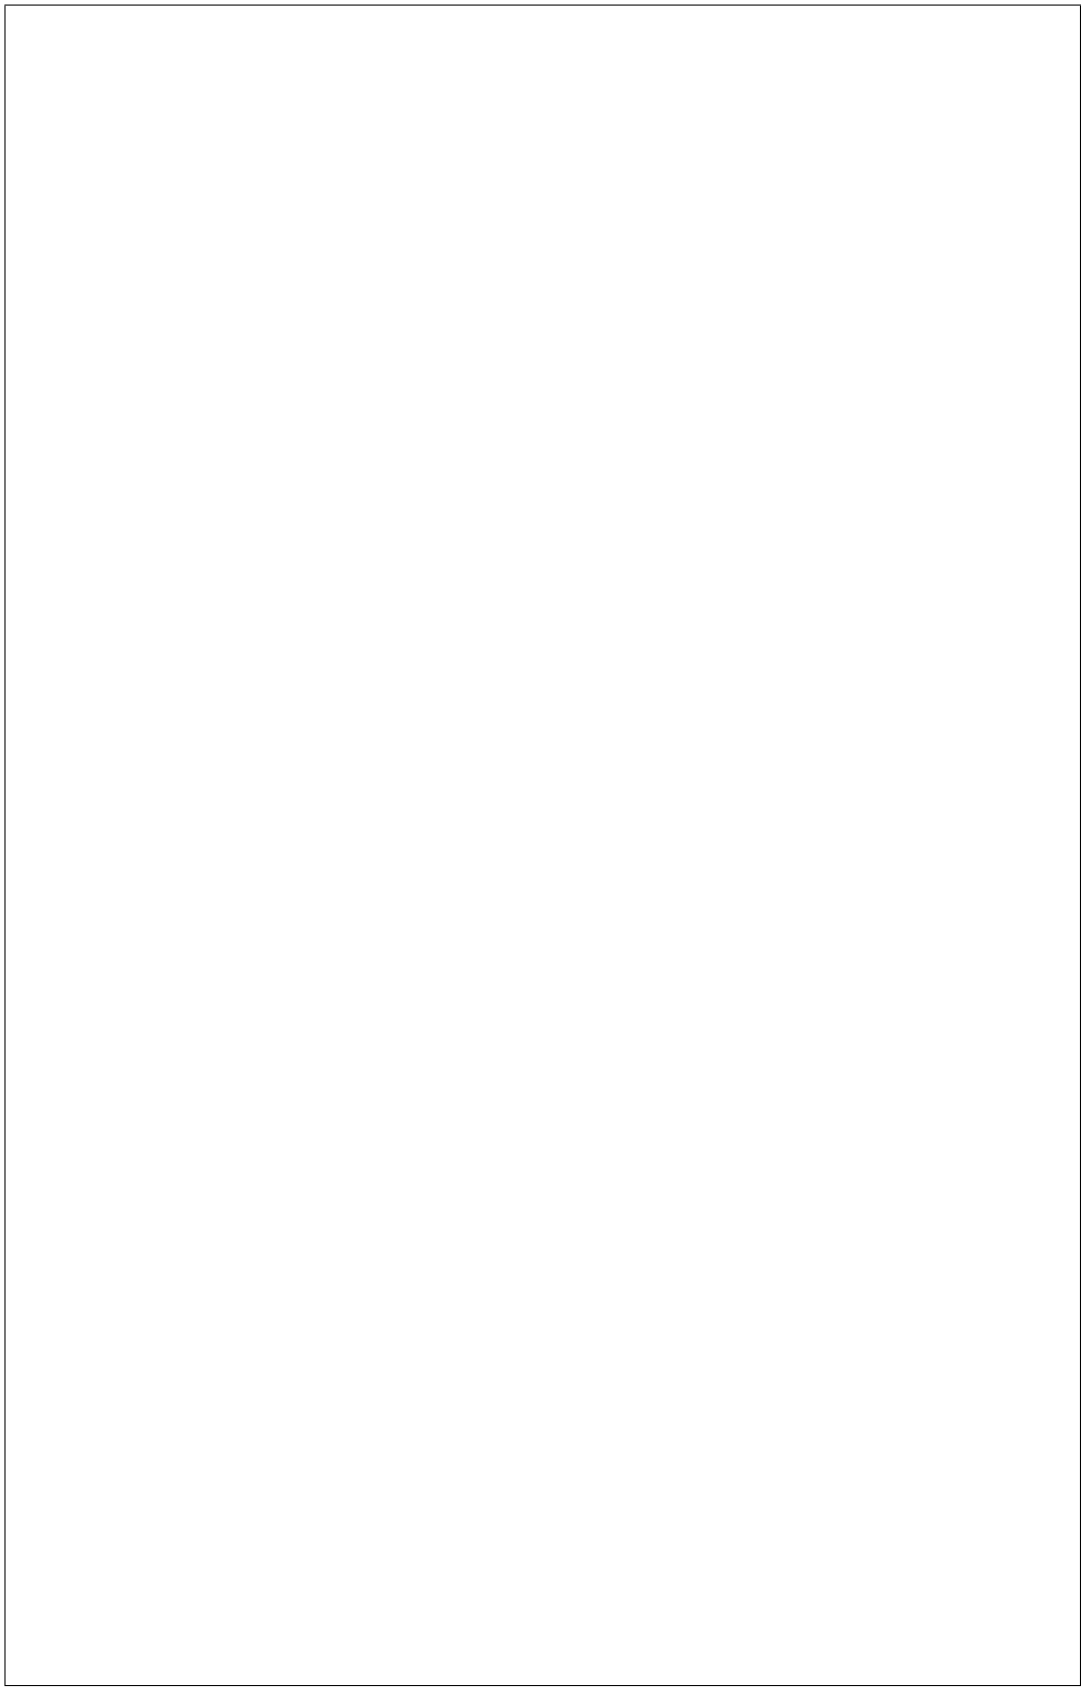

11 Review of applications: volume, area

Main points:

- This section is a review of applications of the double integral such as calculating net volume and area in the plane.
- By this stage you should be comfortable with using a double integral to calculate the net volume below a surface.
- You should know how to find the area of a general region in the plane.

When the regions are more difficult, it is a good idea to draw two diagrams - the 3-D diagram with the x-y-z axes and the 2-D one of the region in the x-y plane.

11.1 Example: Find the volume of the tetrahedron bounded by the planes $x + 2y + z = 2$, $x = 2y$, $x = 0$ and $z = 0$.

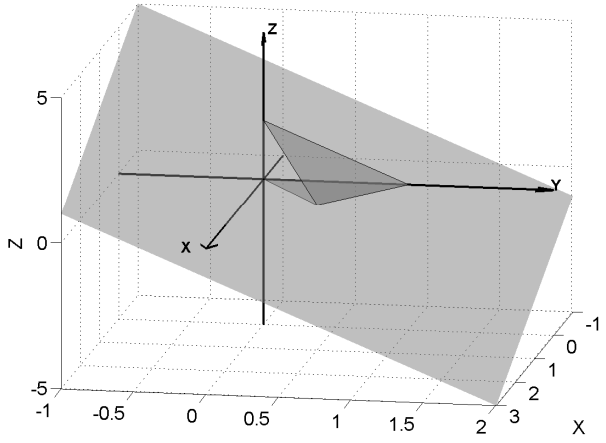

Figure 25: You should be able to reproduce a diagram like this one as an aid to determining the bounds of integration.

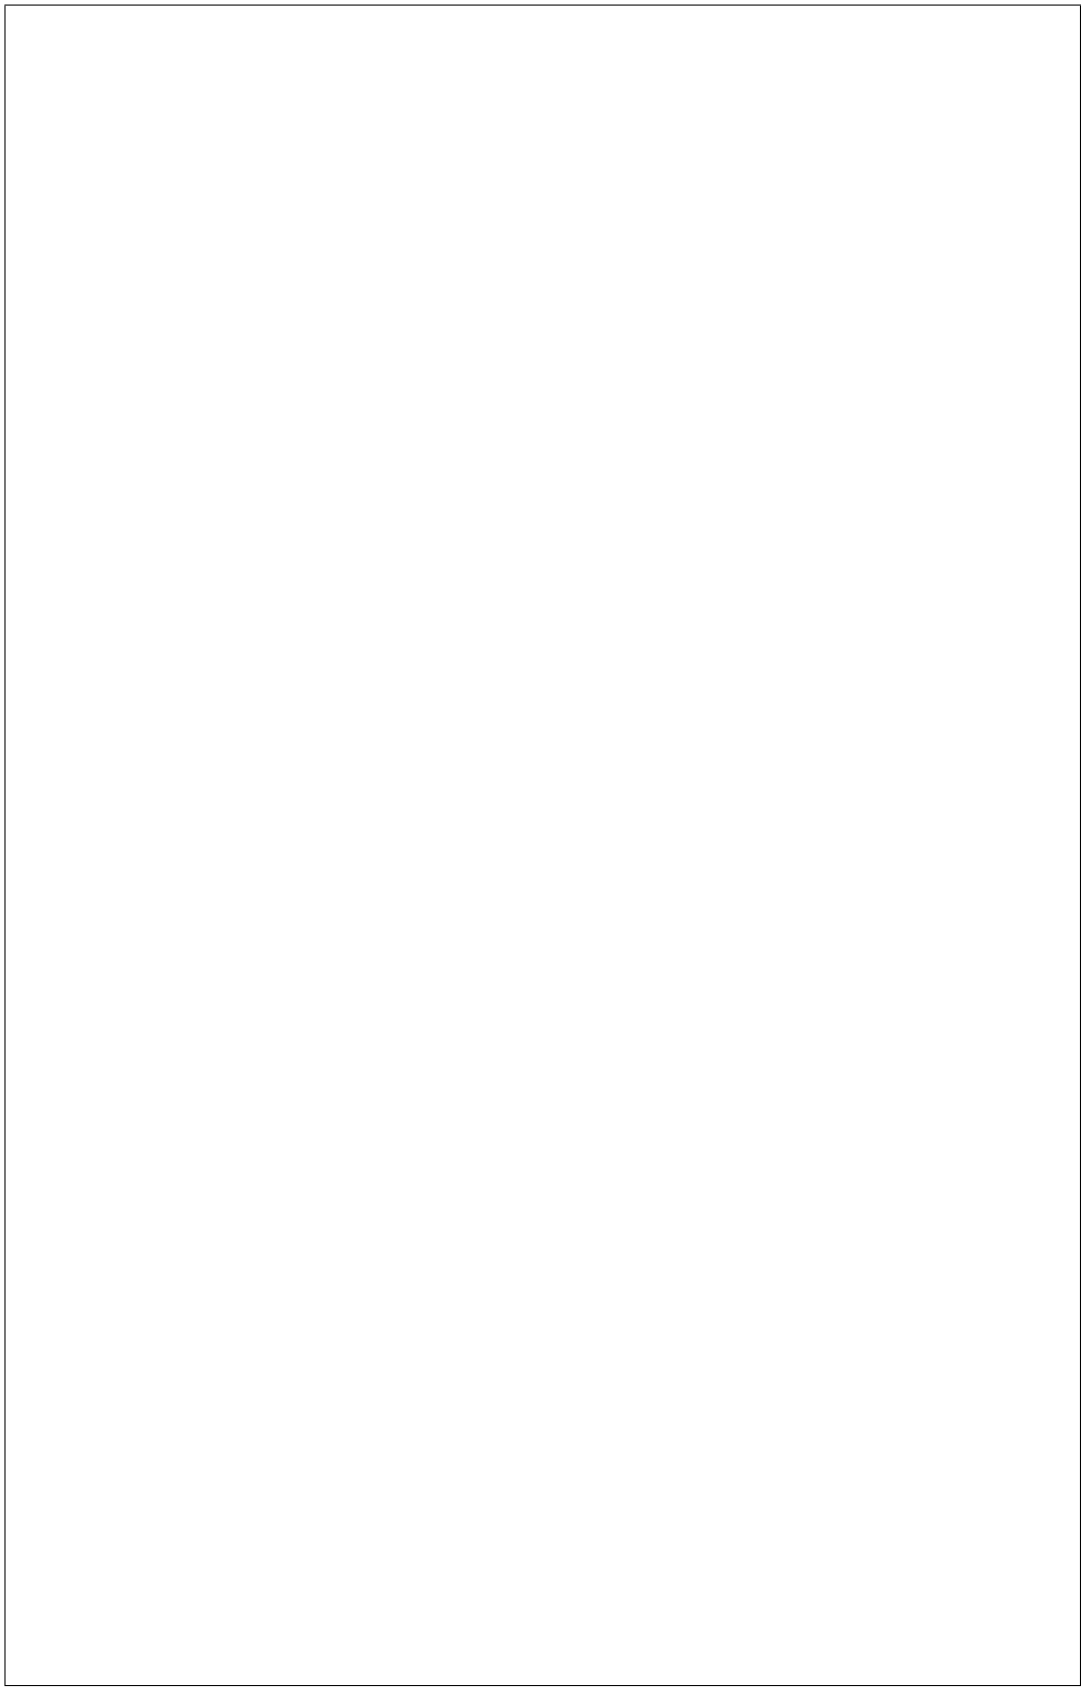

11.2 Area

Note that if we take $f(x, y) = 1$, we have $\iint_D 1 \, dA =$ area of the region D.

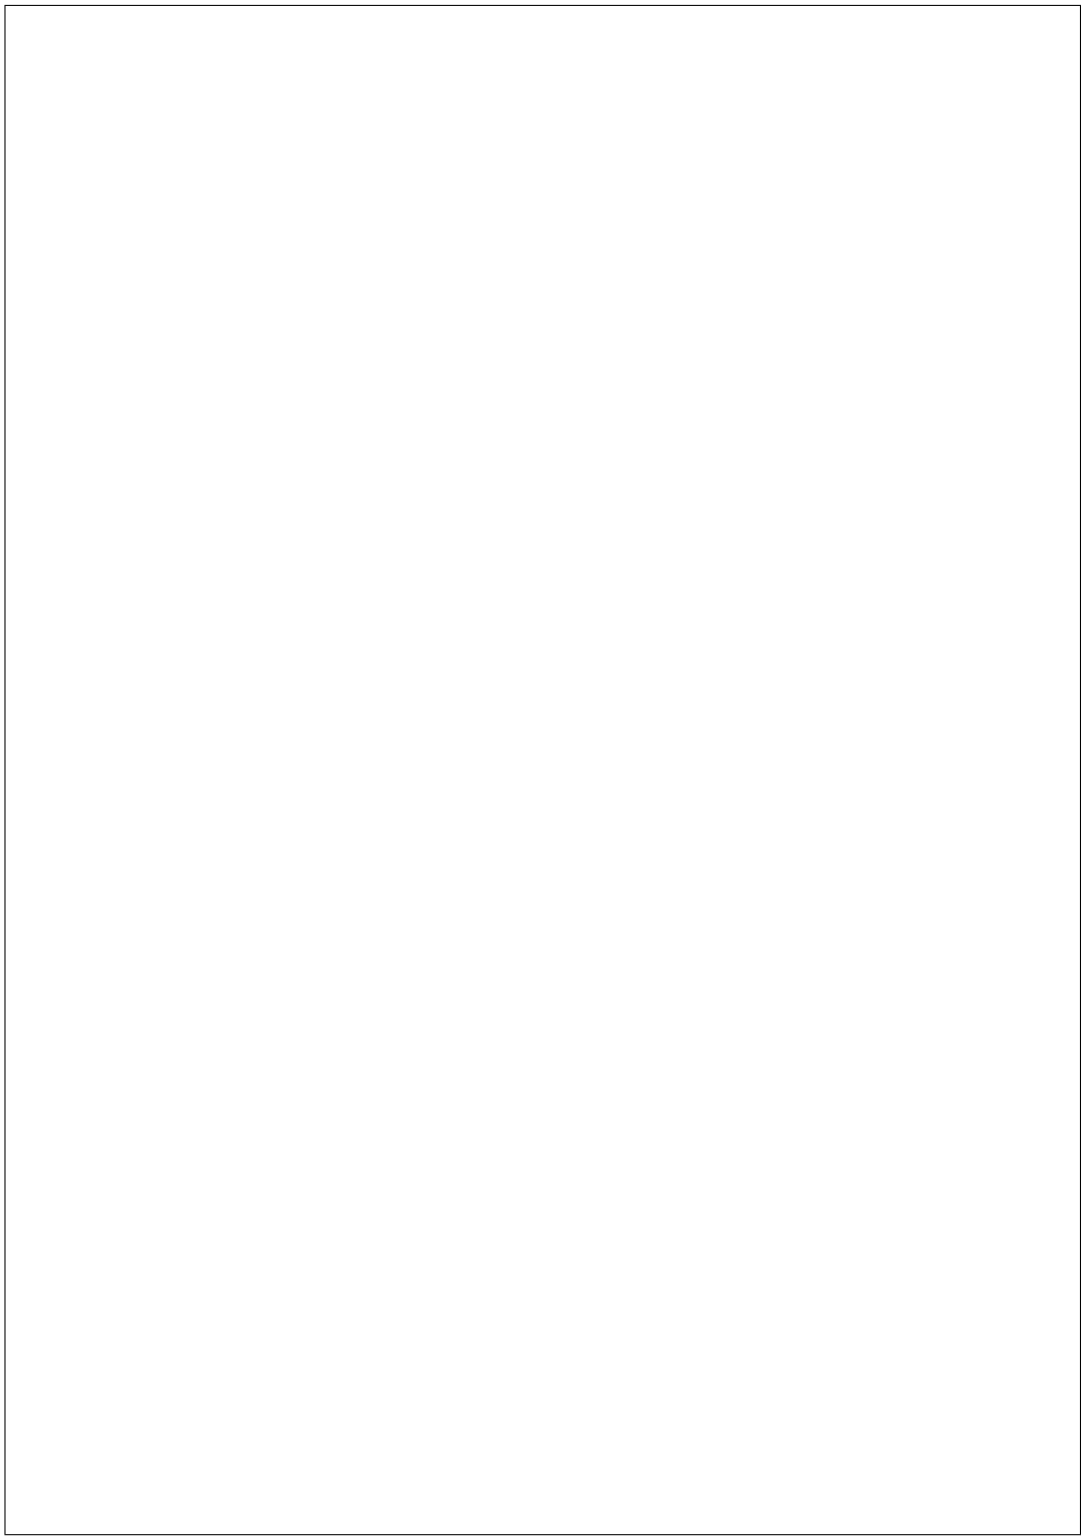

11.3 Find the area enclosed by the ellipse $\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$

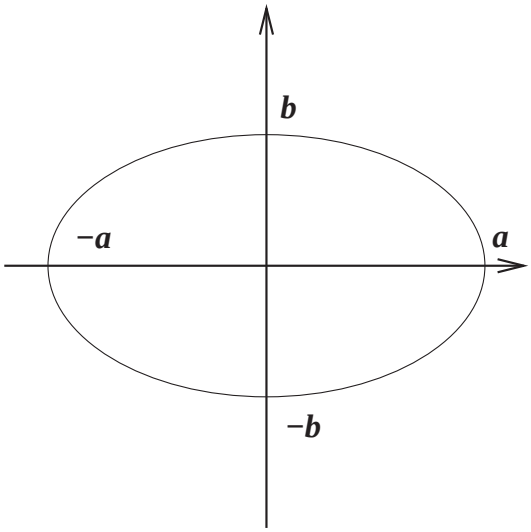

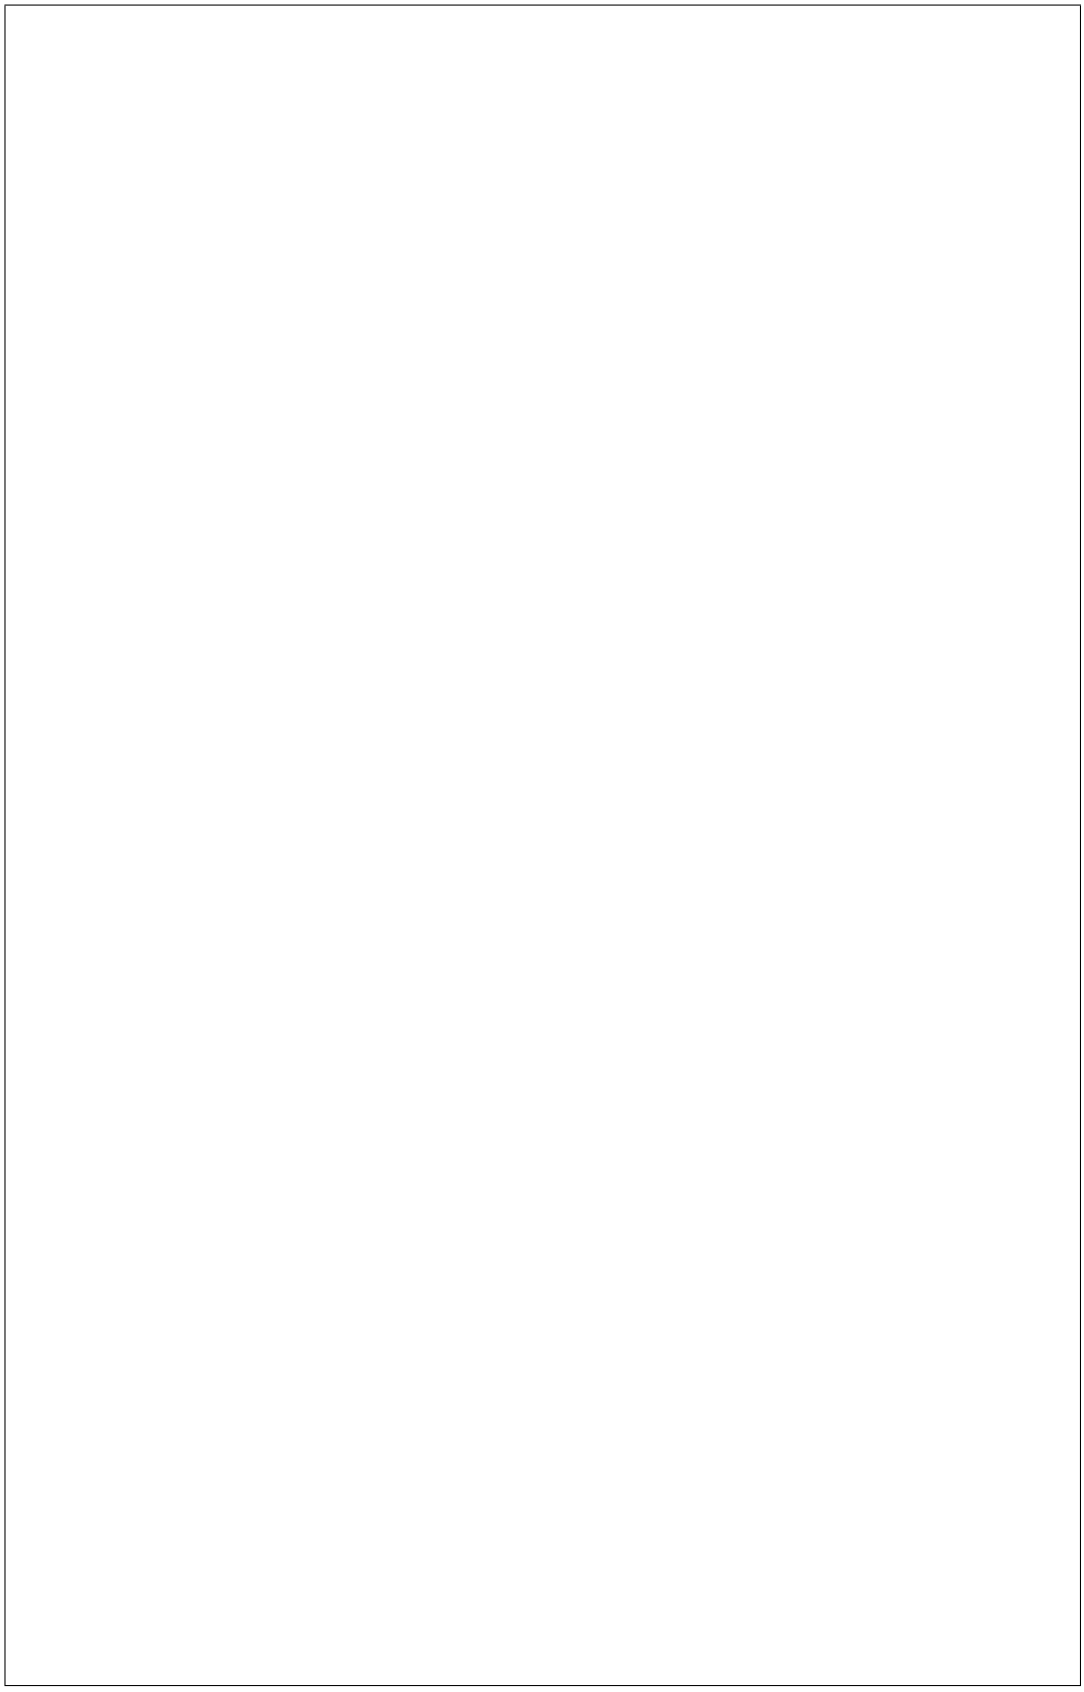

12 Double integrals in polar coordinates

By the end of this section, you should be able to answer the following questions:

- What is the relationship between polar coordinates and rectangular coordinates?
- How do you transform a double integral in rectangular coordinates into one in terms of polar coordinates?
- What is the Jacobian of the transformation?

For annular regions with circular symmetry, rectangular coordinates are difficult. It can be more convenient to use *polar coordinates*.

The following diagram explains the relationship between the polar variables r, θ and the usual rectangular ones x, y.

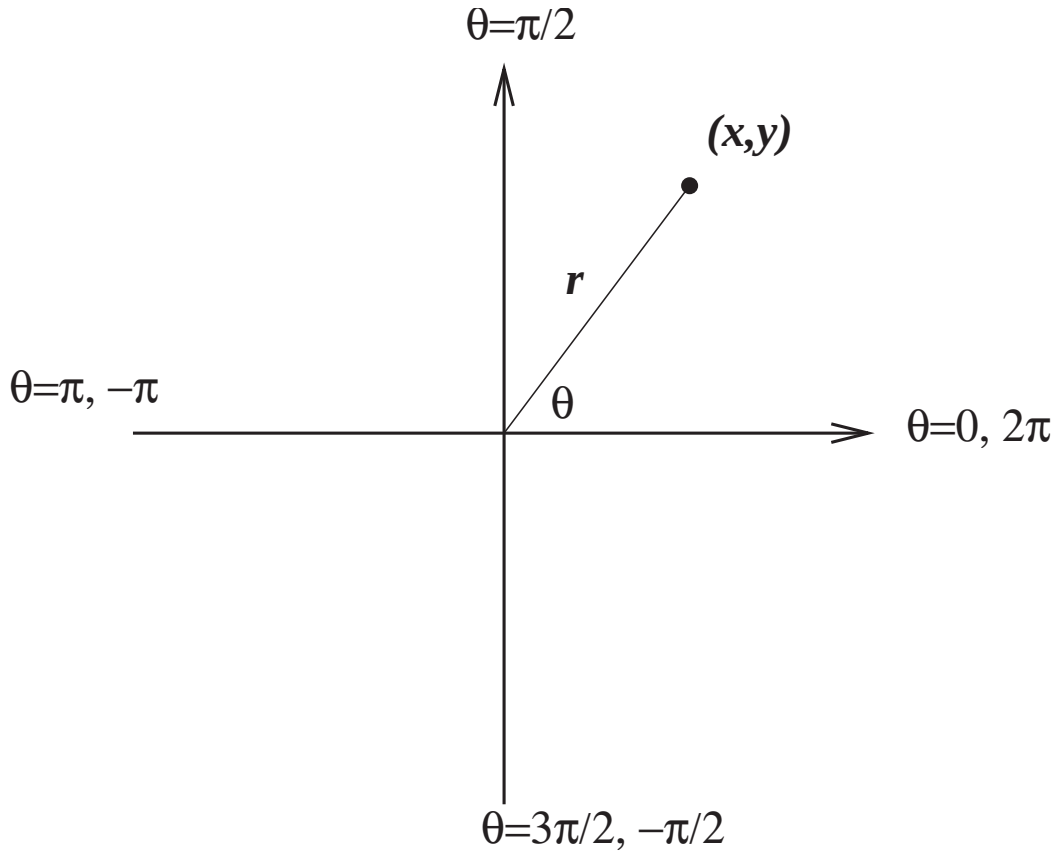

For polar coordinates, we have

$$x = r\cos\theta, \quad y = r\sin\theta.$$

Consider the volume of a solid beneath a surface $z = f(x, y)$ and above a circular region in the x-y plane.

We divide the region into a polar grid as in the following diagram:

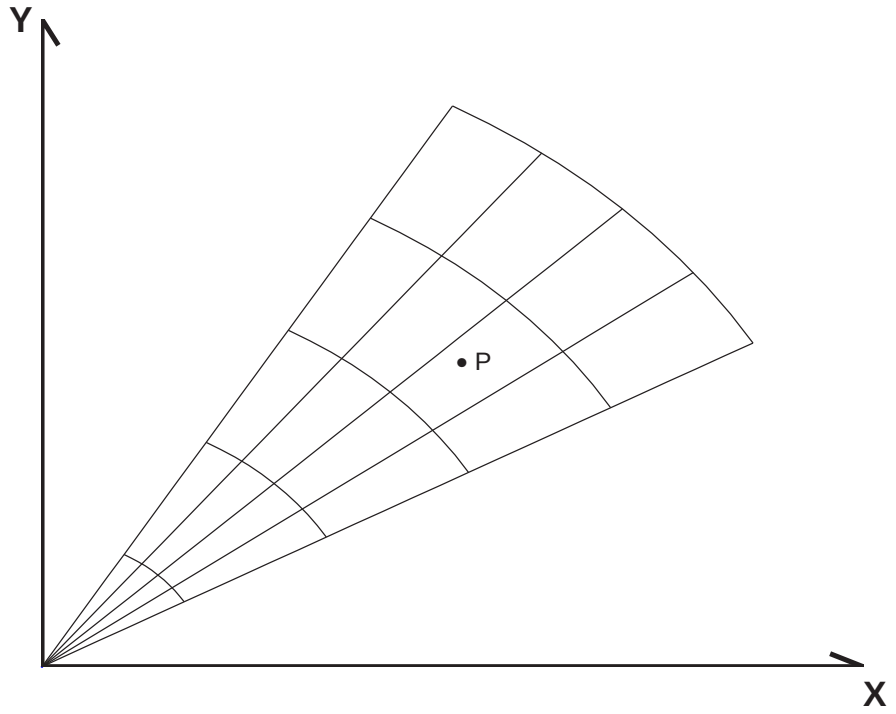

We first approximate the area of each polar rectangle as a regular rectangle. We do this as follows. Choose a point P inside each polar rectangle in the polar grid. Let $P = (x^*, y^*)$ or in polar coordinates $P = (r^*, \theta^*)$, where

$$x^* = r^* \cos\theta^*, \quad y^* = r^* \sin\theta^*.$$

The area of the polar rectangle containing P can be approximated as $r^* \Delta\theta \Delta r$. Therefore the volume under the surface and above each polar rectangle can be approximated as

$$\text{vol. one box } \approx r^* \Delta\theta \Delta r f(r^* \cos\theta^*, r^* \sin\theta^*).$$

Here $f(r^* \cos\theta^*, r^* \sin\theta^*)$ is the value of the function at the point P, which is also the height of the box used in the approximation. To obtain an approximation for

the entire volume below the surface, we sum over the entire polar grid:

$$\begin{aligned} \text{vol.} &\approx \sum_{\text{(polar grid)}} r^*\Delta\theta\Delta r f(r^*\cos\theta^*, r^*\sin\theta^*) \\ \Rightarrow \quad \text{vol.} &= \lim_{\Delta r,\Delta\theta\to 0} \sum_{\text{(polar grid)}} r^*\Delta\theta\Delta r f(r^*\cos\theta^*, r^*\sin\theta^*) \\ &= \iint_D f(r\cos\theta, r\sin\theta) r \; d\theta \; dr. \end{aligned}$$

The double integral in rectangular coordinates is then transformed as follows:

$$\iint_R f(x,y) \; dx \; dy = \iint_S f(r\cos\theta, r\sin\theta) r \; dr \; d\theta.$$

12.1 Example: Find $\iint_D e^{-(x^2+y^2)} \; dx \; dy$ where D is the region bounded by the circle $x^2 + y^2 = R^2$.

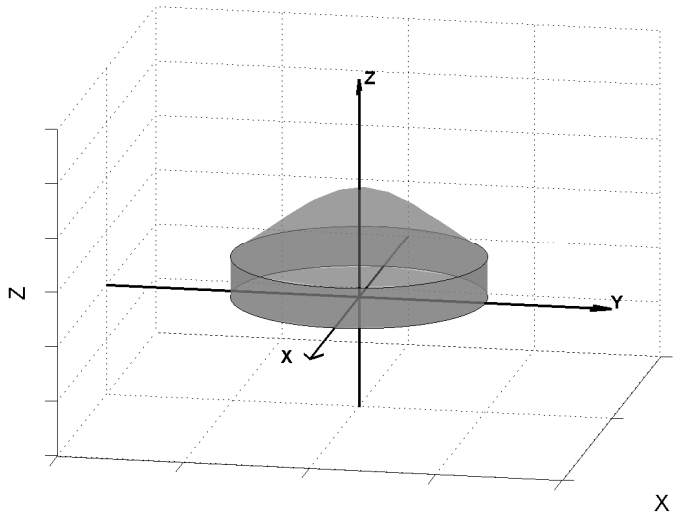

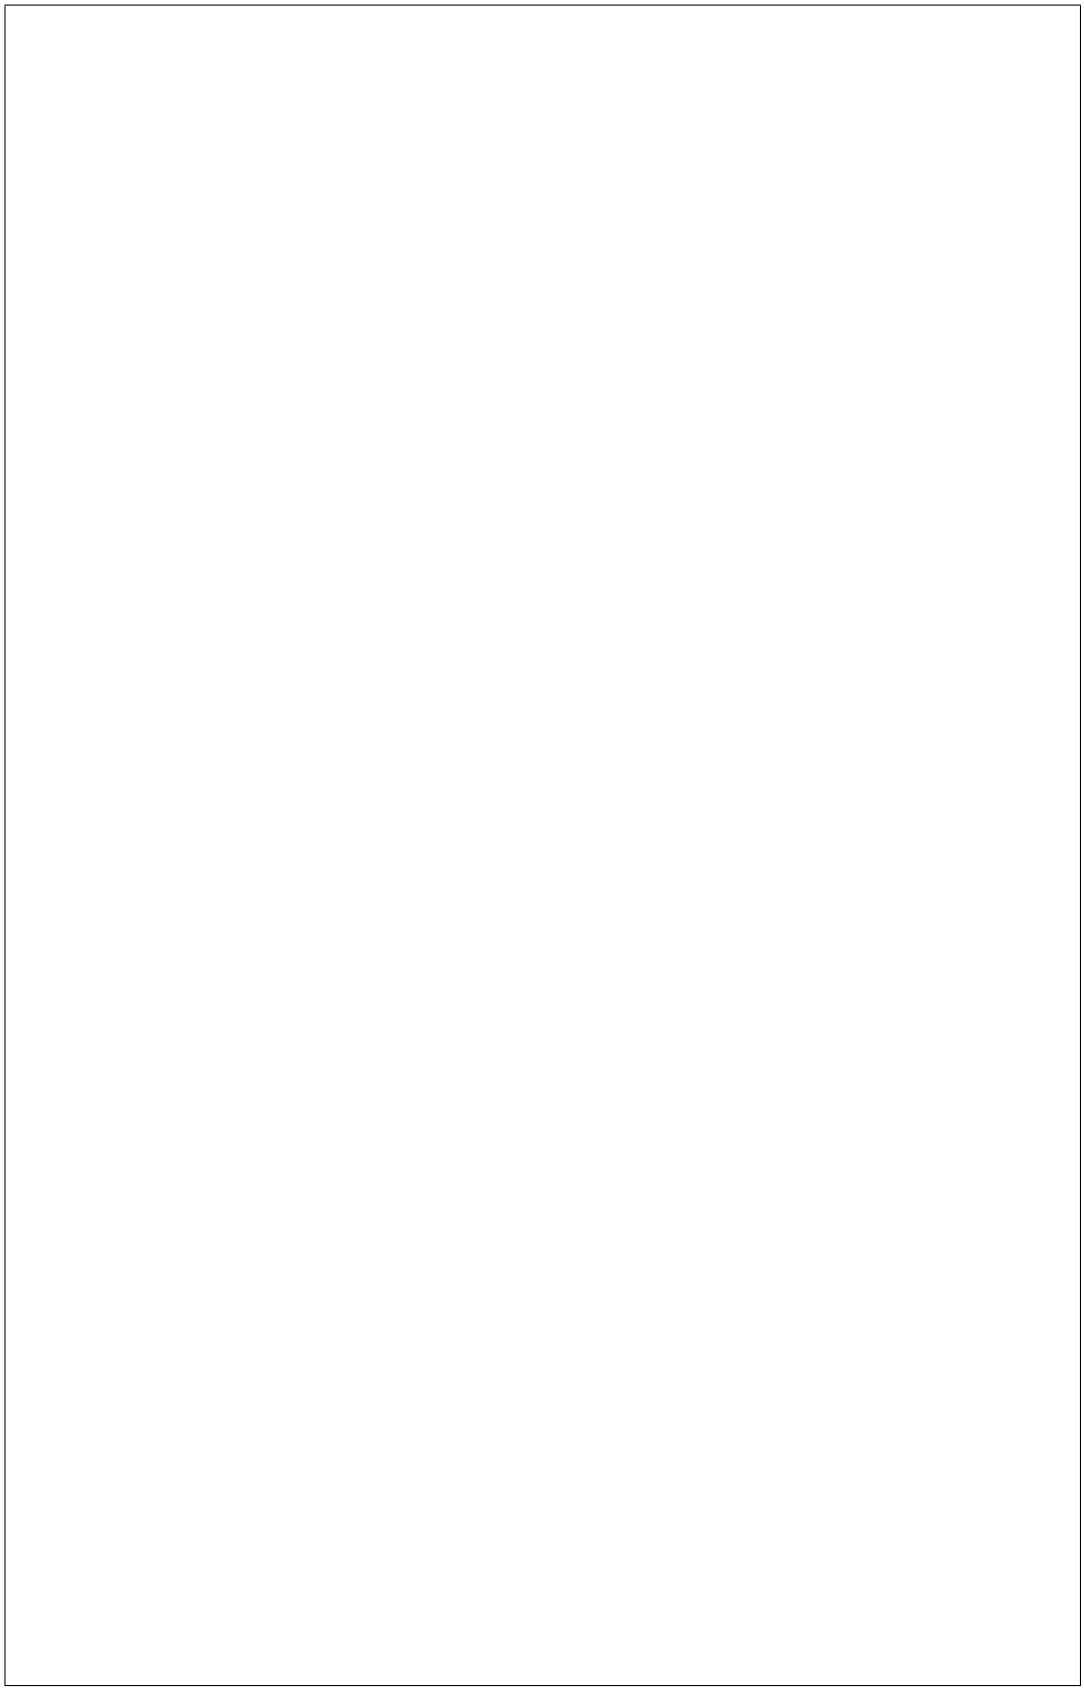

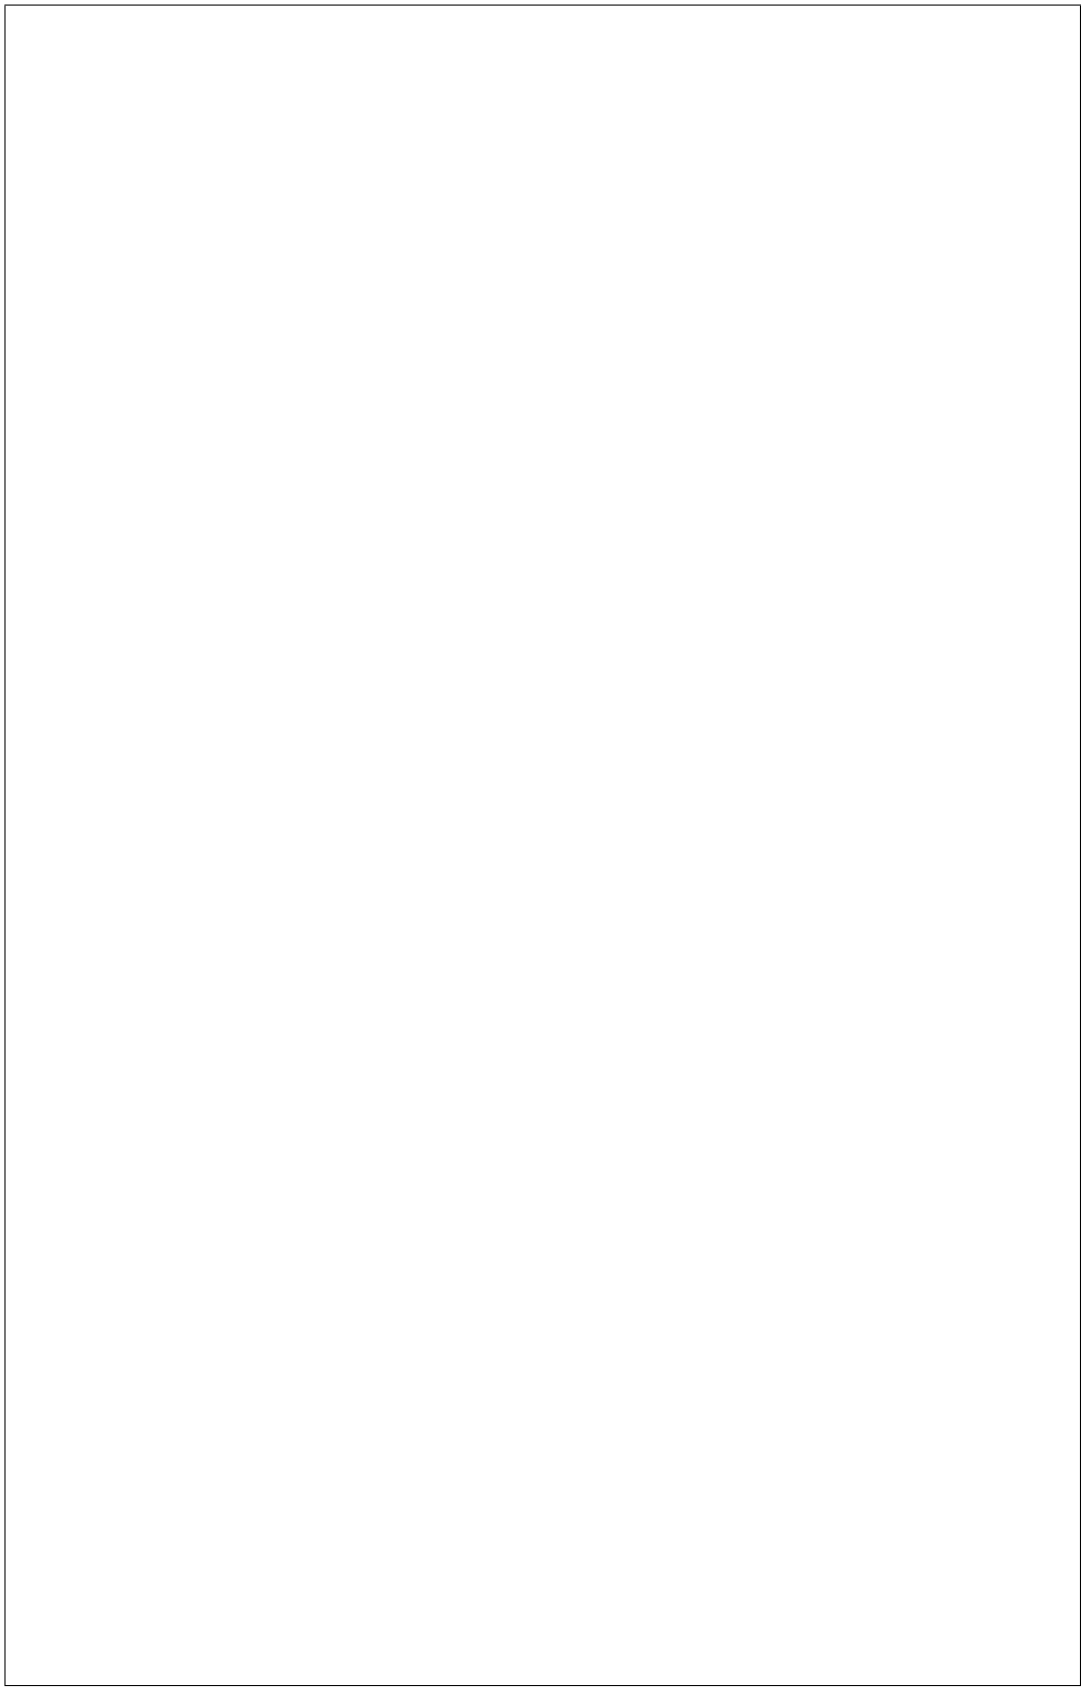

12.2 Example: Find the volume of the solid bounded by the plane $z = 0$ and the paraboloid $z = 1 - x^2 - y^2$.

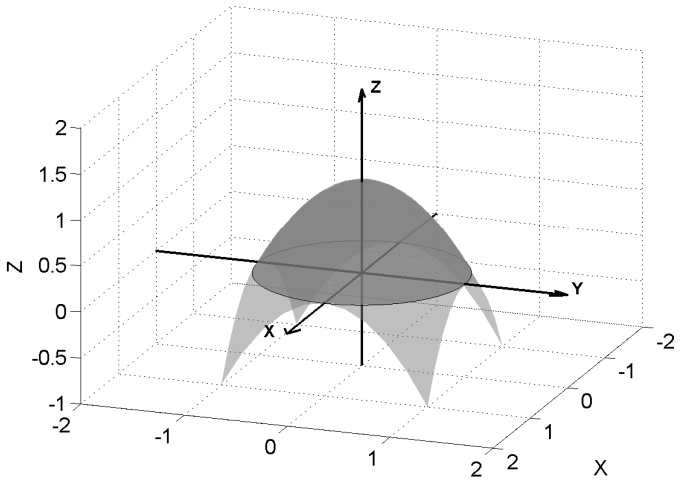

12.3 Find the volume of the solid that lies under the paraboloid $z = x^2 + y^2$ and inside the cylinder $x^2 + y^2 = 2x$, for $z \geq 0$.

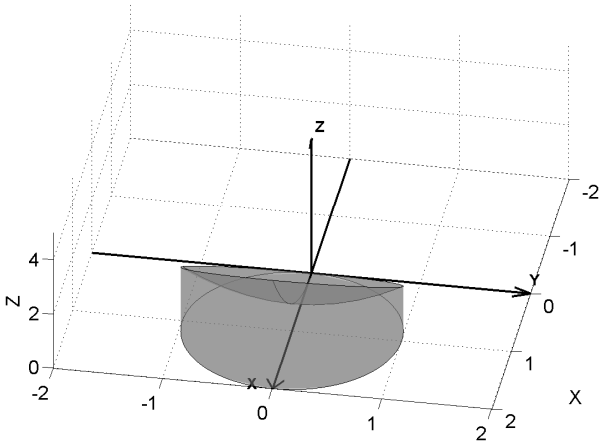

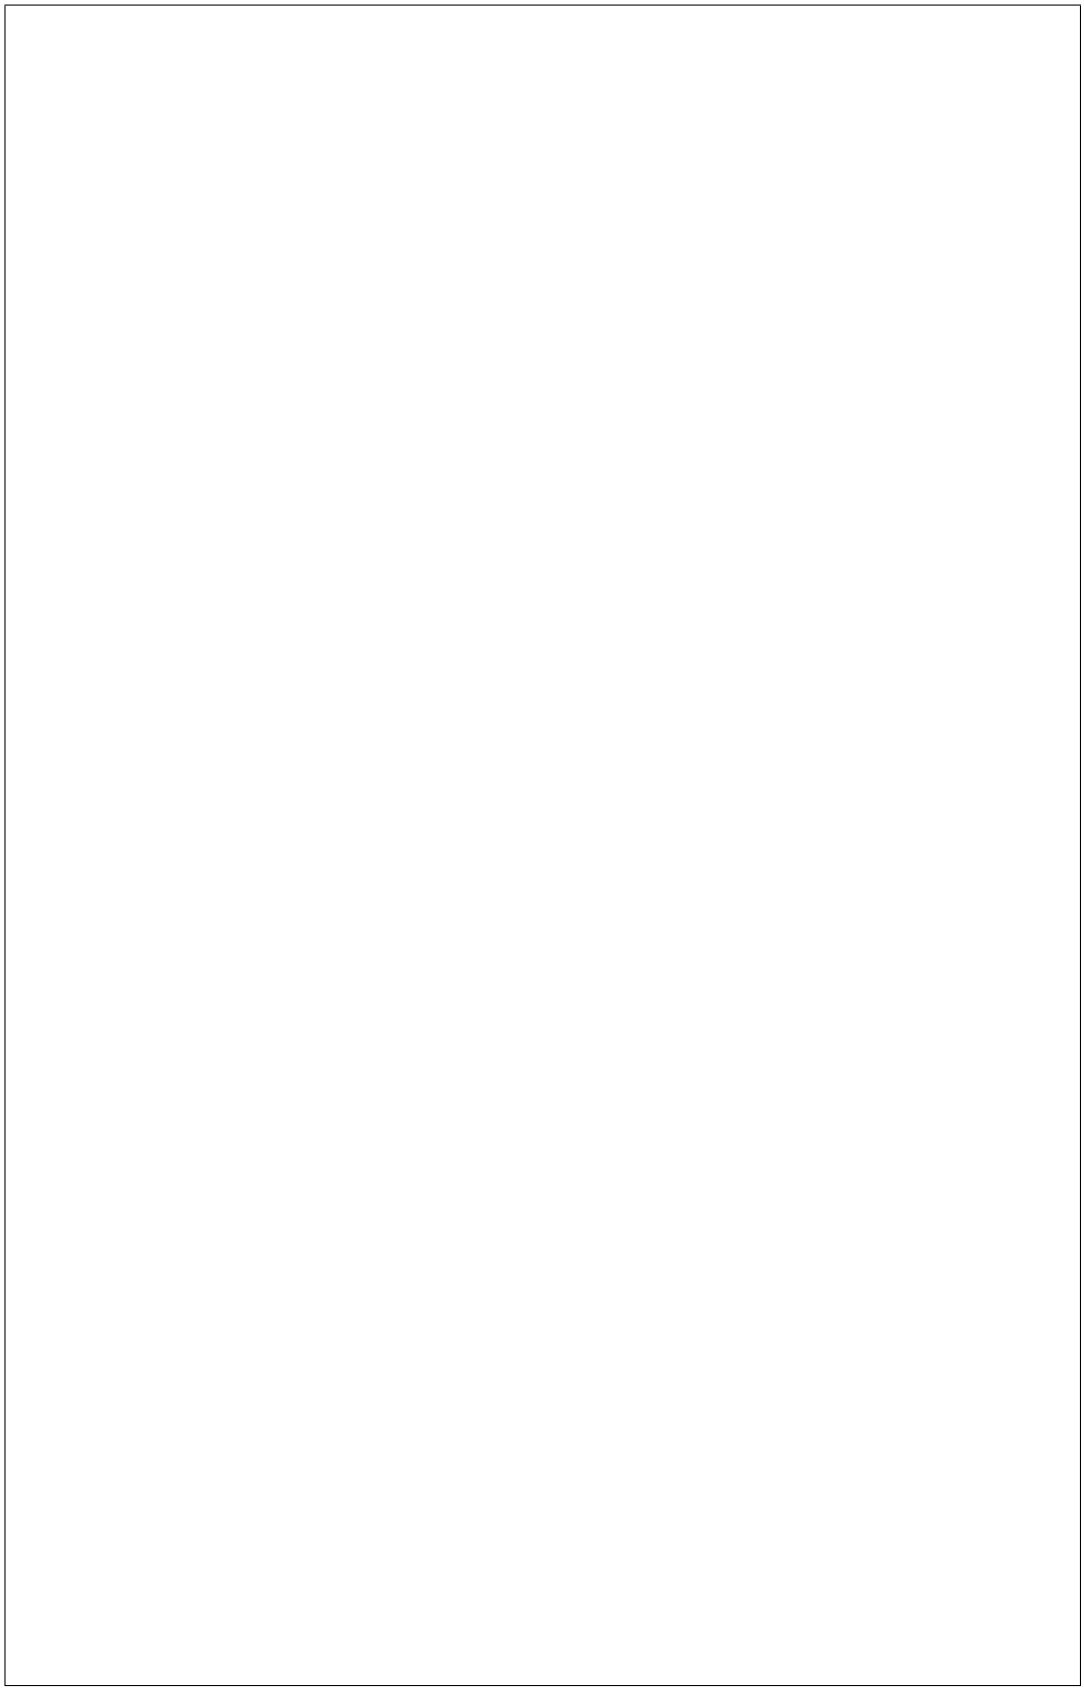

13 Mass, centre of mass and moments

By the end of this section, you should be able to answer the following questions:

- How can we use a double integral to find the mass of a two dimensional object if the density function is known?
- How do we use double integrals to locate the centre of mass of such an object?
- How do we calculate the moments of such an object about the coordinate axes?

Ultimately we want to find a point P on which a thin plate of any given shape balances horizontally. Such a point is called the centre of mass of the plate.

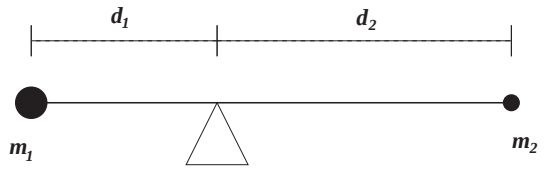

Consider a rod of negligible mass balanced on a fulcrum. The rod has masses m_1 and m_2 at either end, which are a distance d_1 and d_2 respectively from the fulcrum. Because the rod is balanced, we have (thanks to Archimedes) the relationship

$$m_1 d_1 = m_2 d_2.$$

Now suppose the rod lies on the x-axis with m_1 at $x = x_1$, m_2 at $x = x_2$ and the centre of mass at $\overline{x}$.

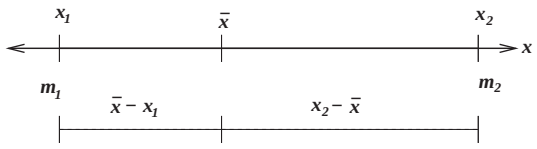

In this case we can write $d_1 = \overline{x} - x_1$ and $d_2 = x_2 - \overline{x}$, so Archimedes' relationship can be expressed

$$m_1(\overline{x} - x_1) = m_2(x_2 - \overline{x}) \;\Rightarrow\; \overline{x} = \frac{m_1 x_1 + m_2 x_2}{m_1 + m_2}.$$

The numbers $m_1 x_1$ and $m_2 x_2$ are called the *moments* of the masses m_1 and m_2 respectively.

In general, a one dimensional system of n "particles" with masses $m_1, \ldots, m_n$ located at $x = x_1, \ldots, x_n$ has its centre of mass located at

$$\overline{x} = \frac{\sum_{i=1}^{n} m_i x_i}{\sum_{i=1}^{n} m_i} = \frac{M}{m}$$

where $m = \sum m_i$ is the total mass of the system and the sum of the individual moments $M = \sum m_i x_i$ is called the moment of the system (with respect to the origin).

Now suppose the rod (which has length l) has mass which is distributed according to the (integrable) density function (mass/unit length)

$$\rho(x) = \lim_{\Delta x \to 0} \frac{\Delta m}{\Delta x}.$$

Consider a small strip of width Δx containing the point x^*. The mass of this strip can be approximated by $\rho(x^*)\Delta x$. Now cut the rod into n strips, and in the same way as above determine (approximately) the mass of each strip. To obtain an approximation for the total mass m of the rod, just add the masses of each n strips:

$$m \approx \sum_{i=1}^{n} \rho(x_i^*)\Delta x_i.$$

To obtain a precise expression for the mass, we take the limit of this sum as $n \to \infty$. In other words,

$$m = \int_0^l \rho(x)dx.$$

We have a similar construction for the moment of the system. Consider the moment of each strip $\approx x_i^* \rho(x_i^*)\Delta x_i$. If we add these, we obtain an approximate expression for the moment of the system:

$$M \approx \sum_{i=1}^{n} x_i^* \rho(x_i^*)\Delta x_i.$$

Taking the limit as $n \to \infty$ we obtain an expression for the moment of the system about the origin:

$$M = \int_0^l x\rho(x)dx.$$

The centre of mass is located at $\overline{x} = M/m$.

Now let's generalize this to two dimensions.

Suppose the lamina occupies a region D in the x-y plane and its density (in units of mass/unit area) is given by an integrable function $\rho(x, y)$. In other words,

$$\rho(x, y) = \lim \frac{\Delta m}{\Delta A},$$

where Δm and ΔA are the mass and area of a small rectangle containing the point (x, y), and the limit is taken as the dimensions of $\Delta A \to 0$.

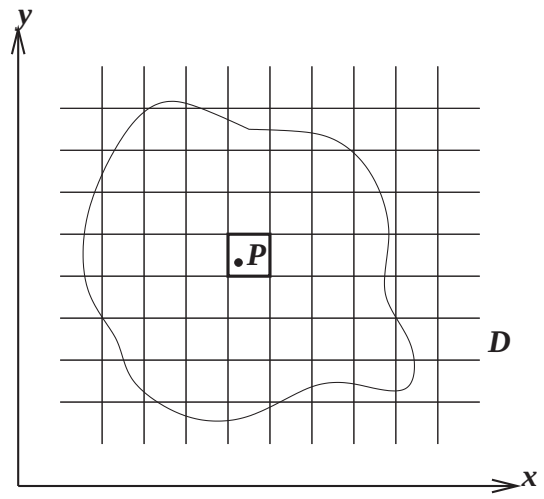

Figure 26: The point $P = (x_i^*, y_j^*)$ in the rectangle R_{ij}.

To approximate the total mass of the lamina, we partition D into small rectangles (say R_{ij}) and choose a point (x_i^*, y_j^*) inside R_{ij}. The mass of the lamina inside R_{ij} is approximately $\rho(x_i^*, y_j^*)\Delta A_{ij}$, where ΔA_{ij} is the area of R_{ij}. Adding all such masses, we have the approximation

$$m \approx \sum_{i=1}^{m} \sum_{j=1}^{n} \rho(x_i^*, y_j^*)\Delta A_{ij}.$$

If we then take the limit as $m, n \to 0$, we obtain

$$m = \iint_D \rho(x, y) dA.$$

In a similar way, we can determine the moment of the lamina about the x-axis to be

$$M_x = \iint_D y\rho(x,y)dA$$

and the moment of the lamina about the y-axis to be

$$M_y = \iint_D x\rho(x,y)dA.$$

The centre of mass is located at coordinates $(\overline{x}, \overline{y})$, where

$$\overline{x} = \frac{M_y}{m}, \quad \overline{y} = \frac{M_x}{m}.$$

13.1 Example: find the centre of mass of a triangular lamina with vertices $(0,0), (1,0)$ and $(0,2)$ with constant density ρ_0.

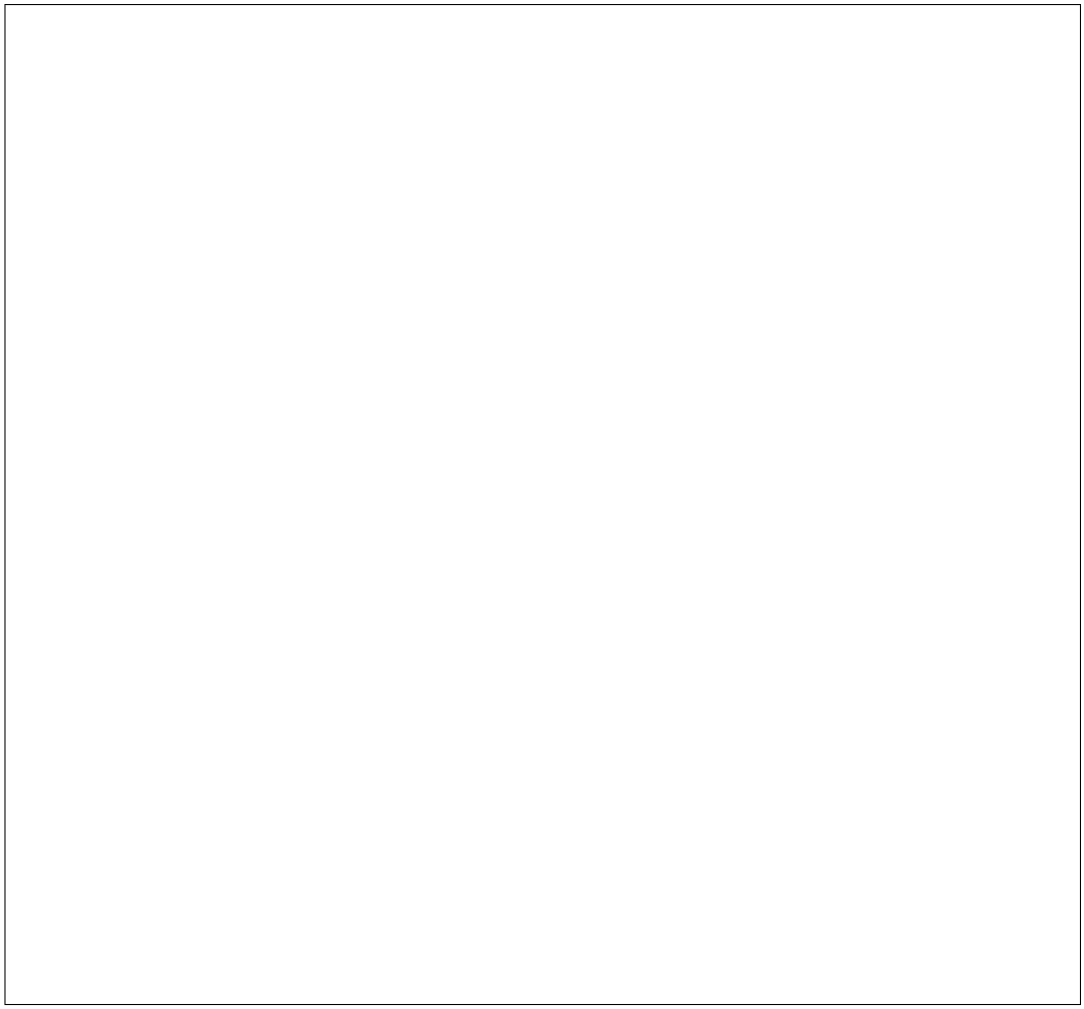

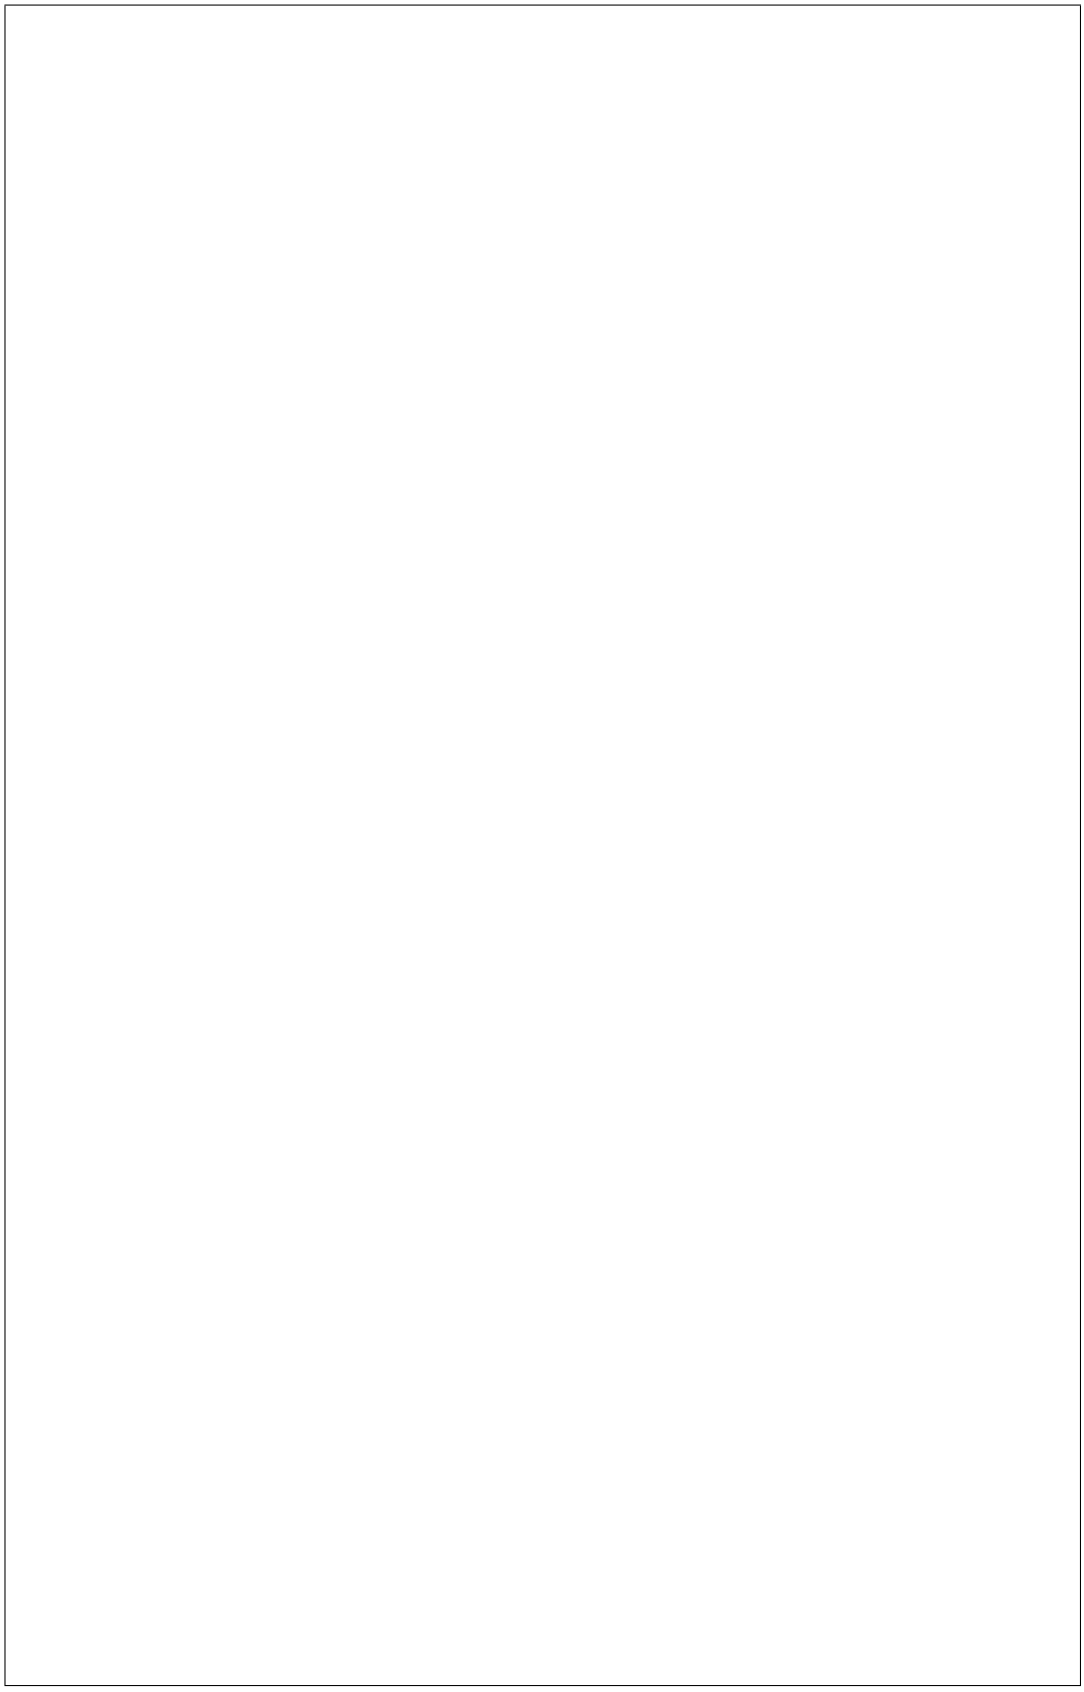

13.2 Example: find the centre of mass of a rectangle with vertices $(0,0),(2,0)$, $(2,1)$ and $(0,1)$ with density $\rho(x,y) = 6x + 12y$.

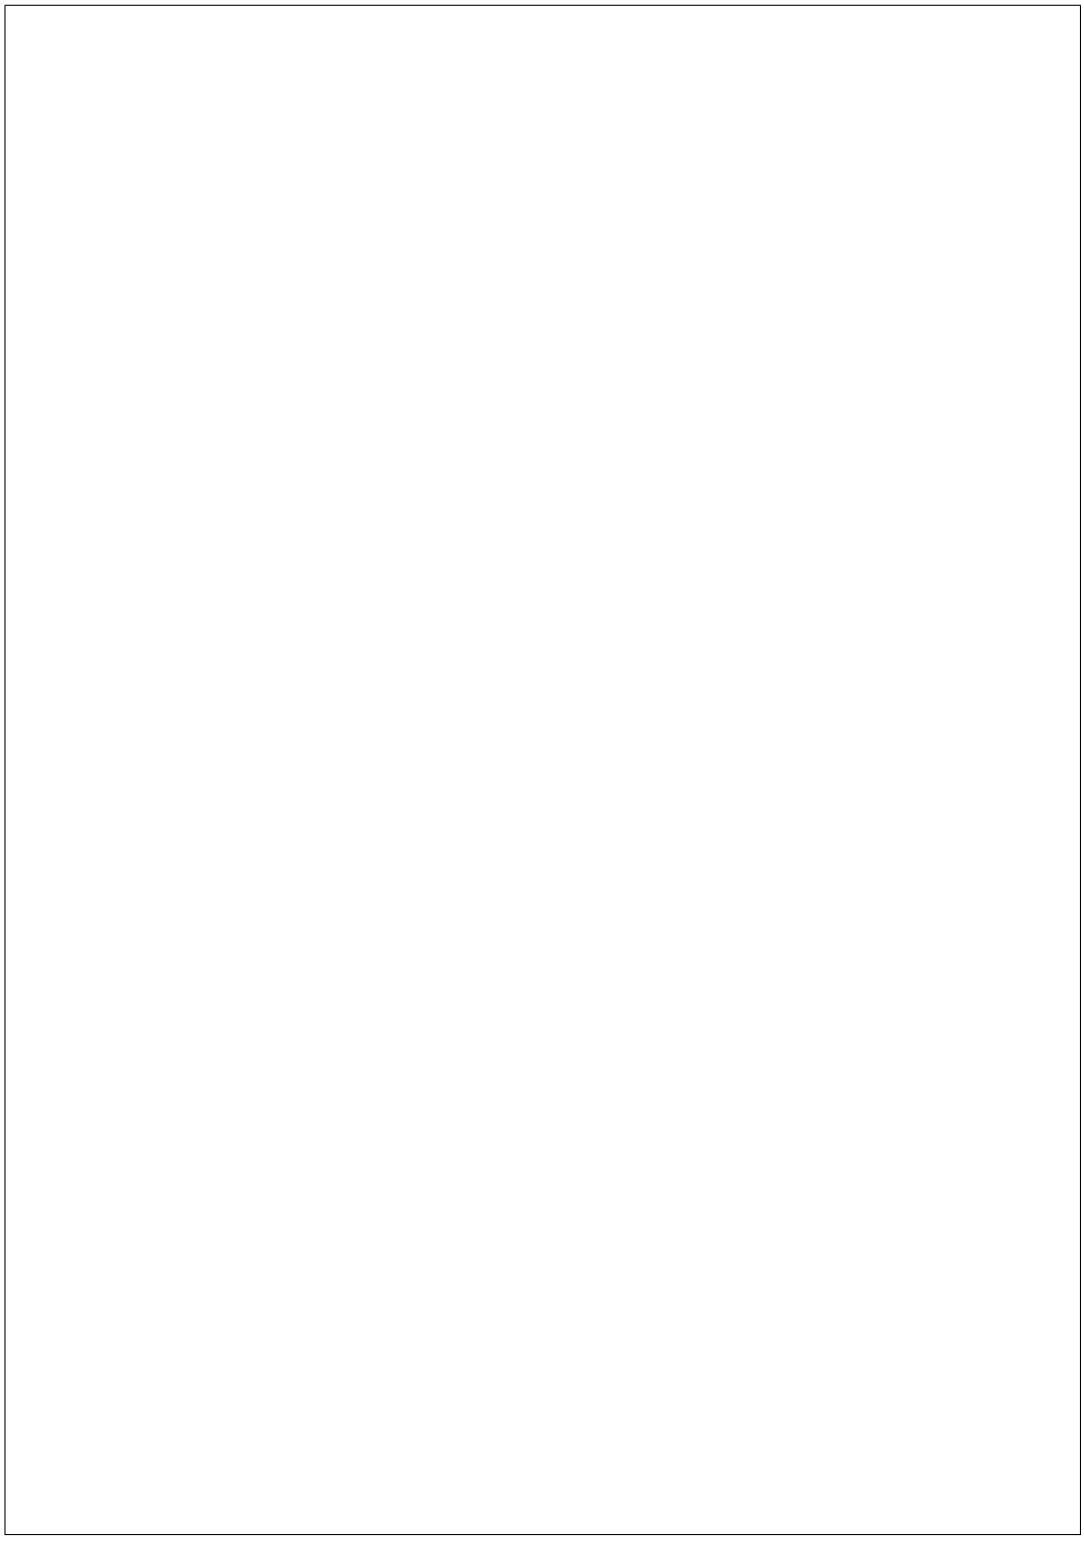

14 Introduction to Triple integrals

By the end of this section, you should be able to answer the following questions:

- How do you evaluate a triple integral?
- How do you use a triple integral to find the mass of a solid object with known density?
- How do you change the order of integration in a triple integral?

We can extend the definition of a double integral to a triple integral

$$\iiint_R f(x,y,z)dV,$$

where R is a region in $\mathbb{R}^3$ and dV is an element of volume.

If R is a region in $\mathbb{R}^3$ specified by

$$\begin{array}{rll} r(x,y) & \leq z \leq & s(x,y) \\ p(x) & \leq y \leq & q(x) \\ a & \leq x \leq & b \end{array} \tag{11}$$

then

$$\iiint_R f(x,y,z)dV$$

$$= \int_a^b \left\{ \int_{p(x)}^{q(x)} \left[\int_{r(x,y)}^{s(x,y)} f(x,y,z)dz \right] dy \right\} dx.$$

In two dimensions, there are 2 possible orders of integration. In three dimensions, there are 6.

14.1 Find the mass of a rectangular block with dimensions $0 \leq x \leq L$, $0 \leq y \leq W$ and $0 \leq z \leq H$ if the density is $\rho = \rho_0 + \alpha xyz$.

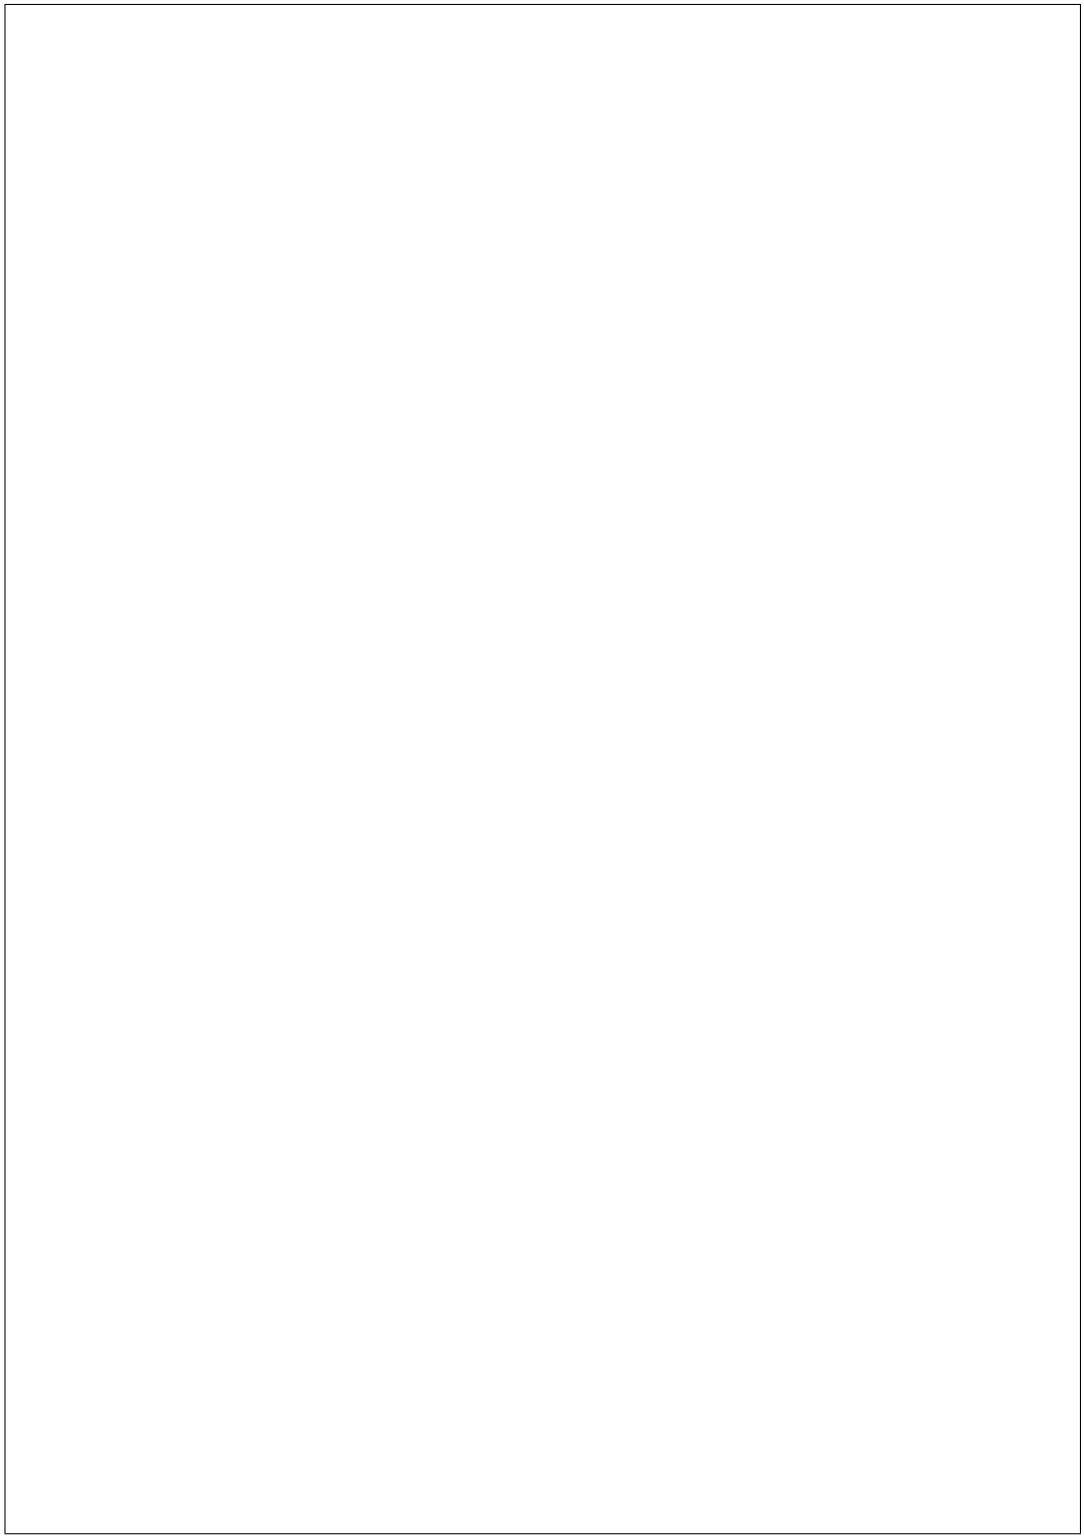

14.2 Evaluate $\iiint_R z \, dV$ over the region R bounded by the surfaces $x = 0$, $y = 0$, $z = 0$ and $x + y + z = 1$.

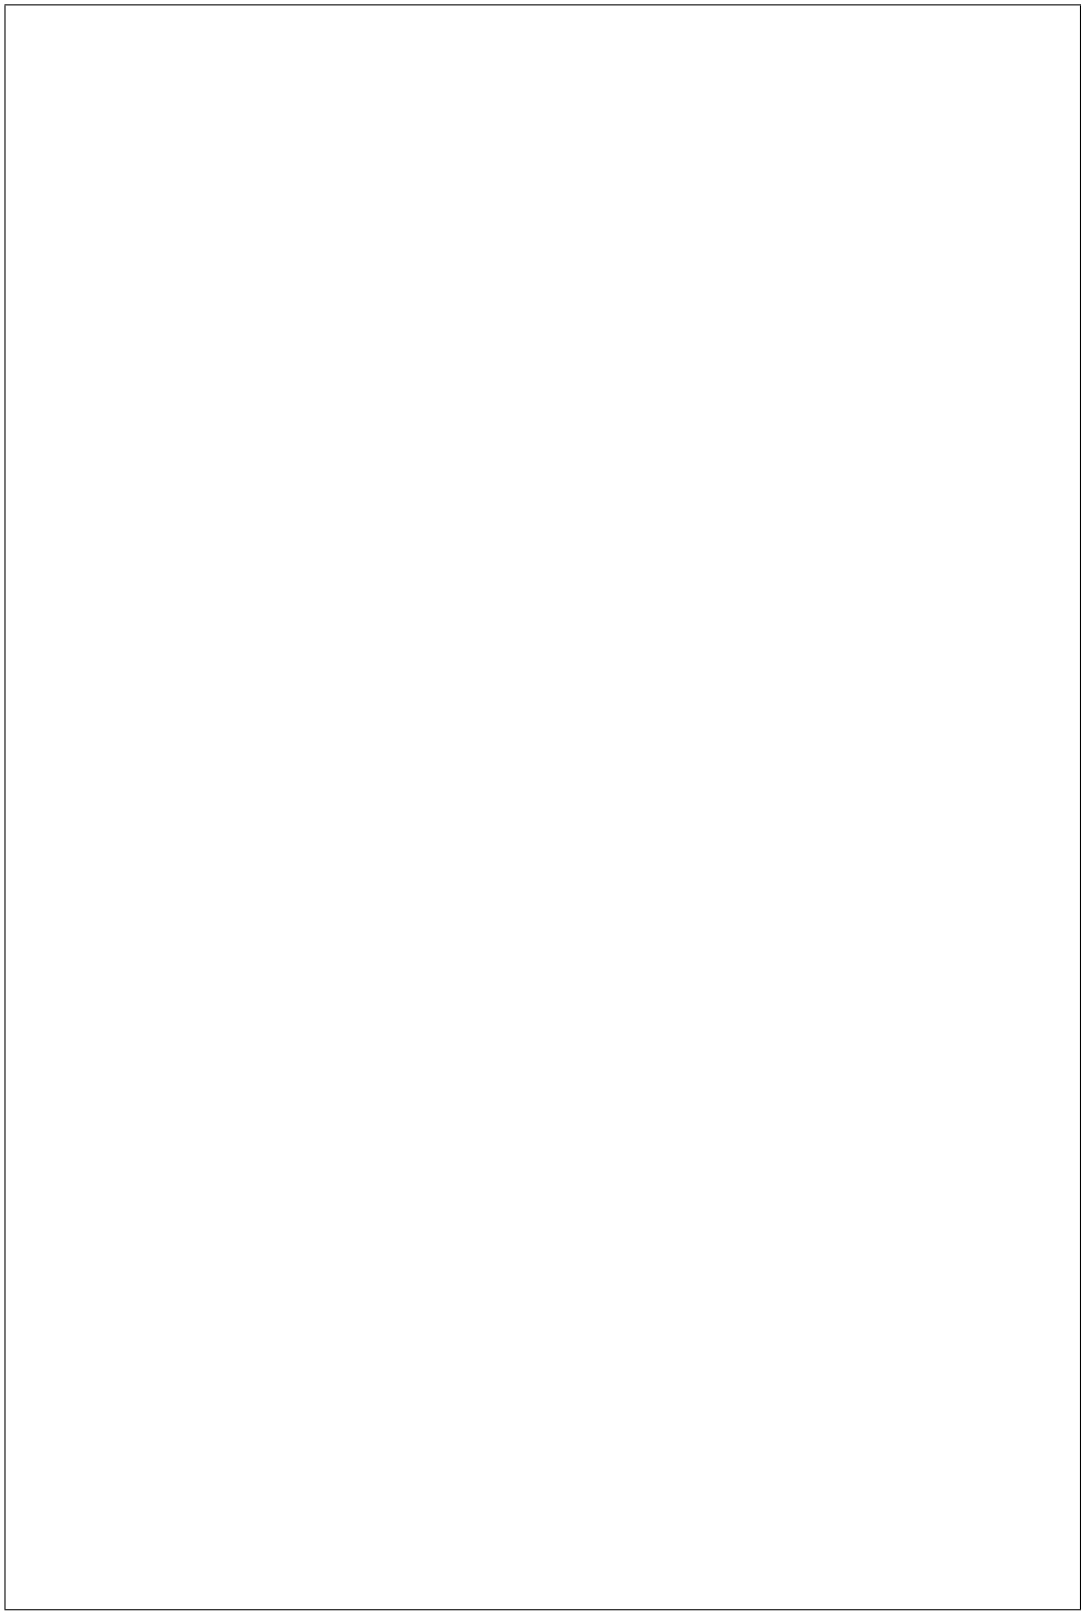

14.3 Changing the order of integration

Express the integral $\int_0^1 \int_{\sqrt{x}}^1 \int_0^{1-y} f(x,y,z)\ dz\ dy\ dx$, in the orders $dz\ dx\ dy$ and $dy\ dz\ dx$.

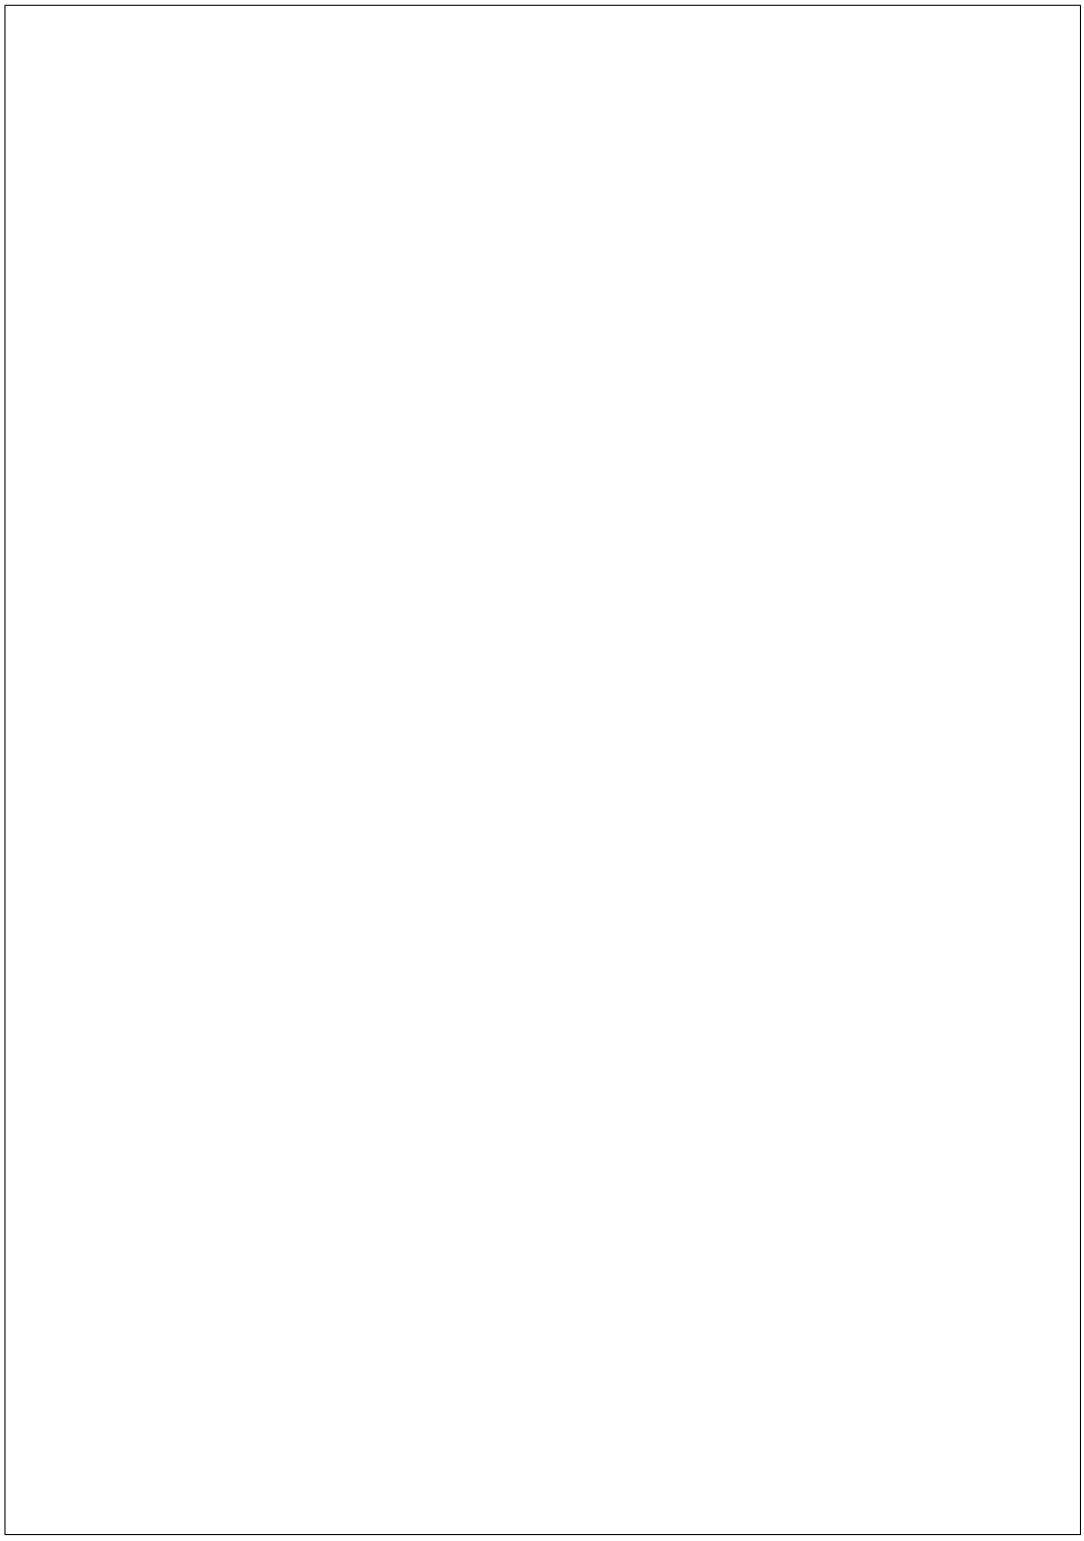

15 Cylindrical coordinates

By the end of this section, you should be able to answer the following questions:

- What is the relationship between rectangular coordinates and cylindrical coordinates?
- How do you transform a triple integral in rectangular coordinates into one in terms of cylindrical coordinates?
- What is the Jacobian of the transformation?

Sometimes it is useful to use cylindrical coordinates in order to simplify the integral. This involves the transformation

$$x = r\cos\theta, \quad y = r\sin\theta, \quad z = z. \tag{12}$$

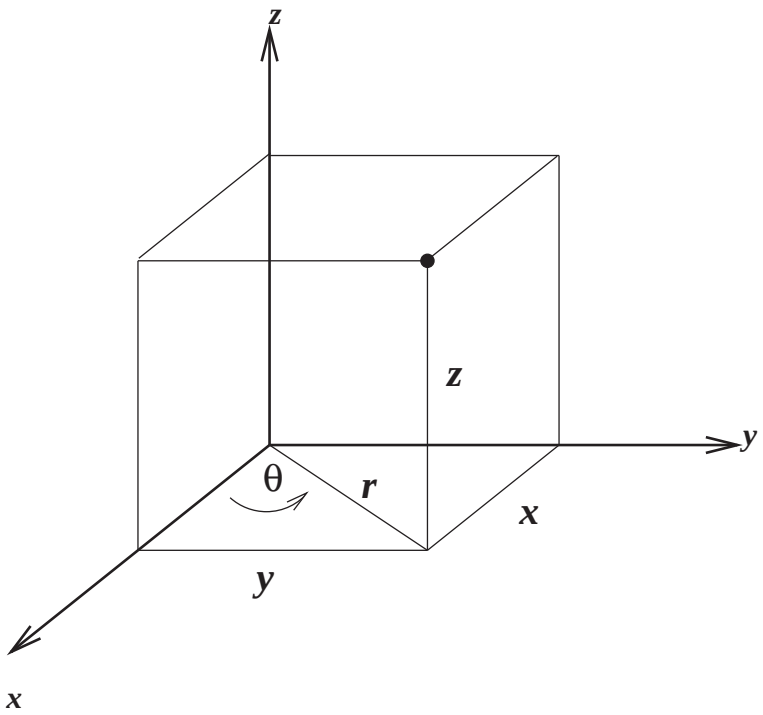

We now aim to calculate a small element of volume of a cylindrical shell. This will then show how in a triple integral we can transform from rectangular coordinates to cylindrical coordinates by substituting the transformation (12) and by making the change

$$dx\ dy\ dz \quad \longrightarrow \quad r\ dr\ d\theta\ dz.$$

Consider the following diagram.

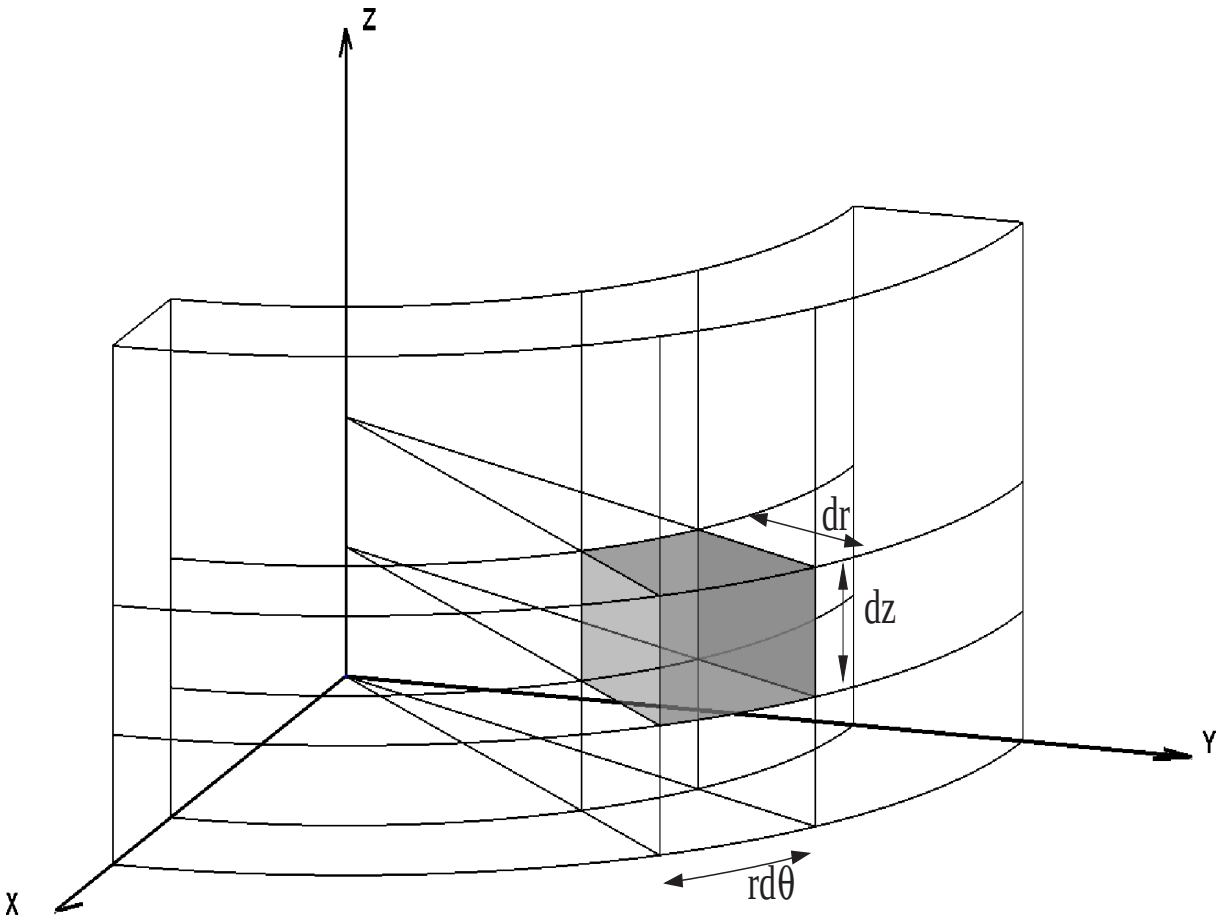

The important result is that the triple integral in rectangular coordinates transforms as follows:

$$\iiint_R f(x,y,z)\ dx\ dy\ dz = \iiint_C f(r\cos\theta, r\sin\theta, z)\ r\ dr\ d\theta\ dz.$$

15.1 A simple example: Find the volume of a cylinder of radius R and height H. (Ans. $\pi R^2 H$)

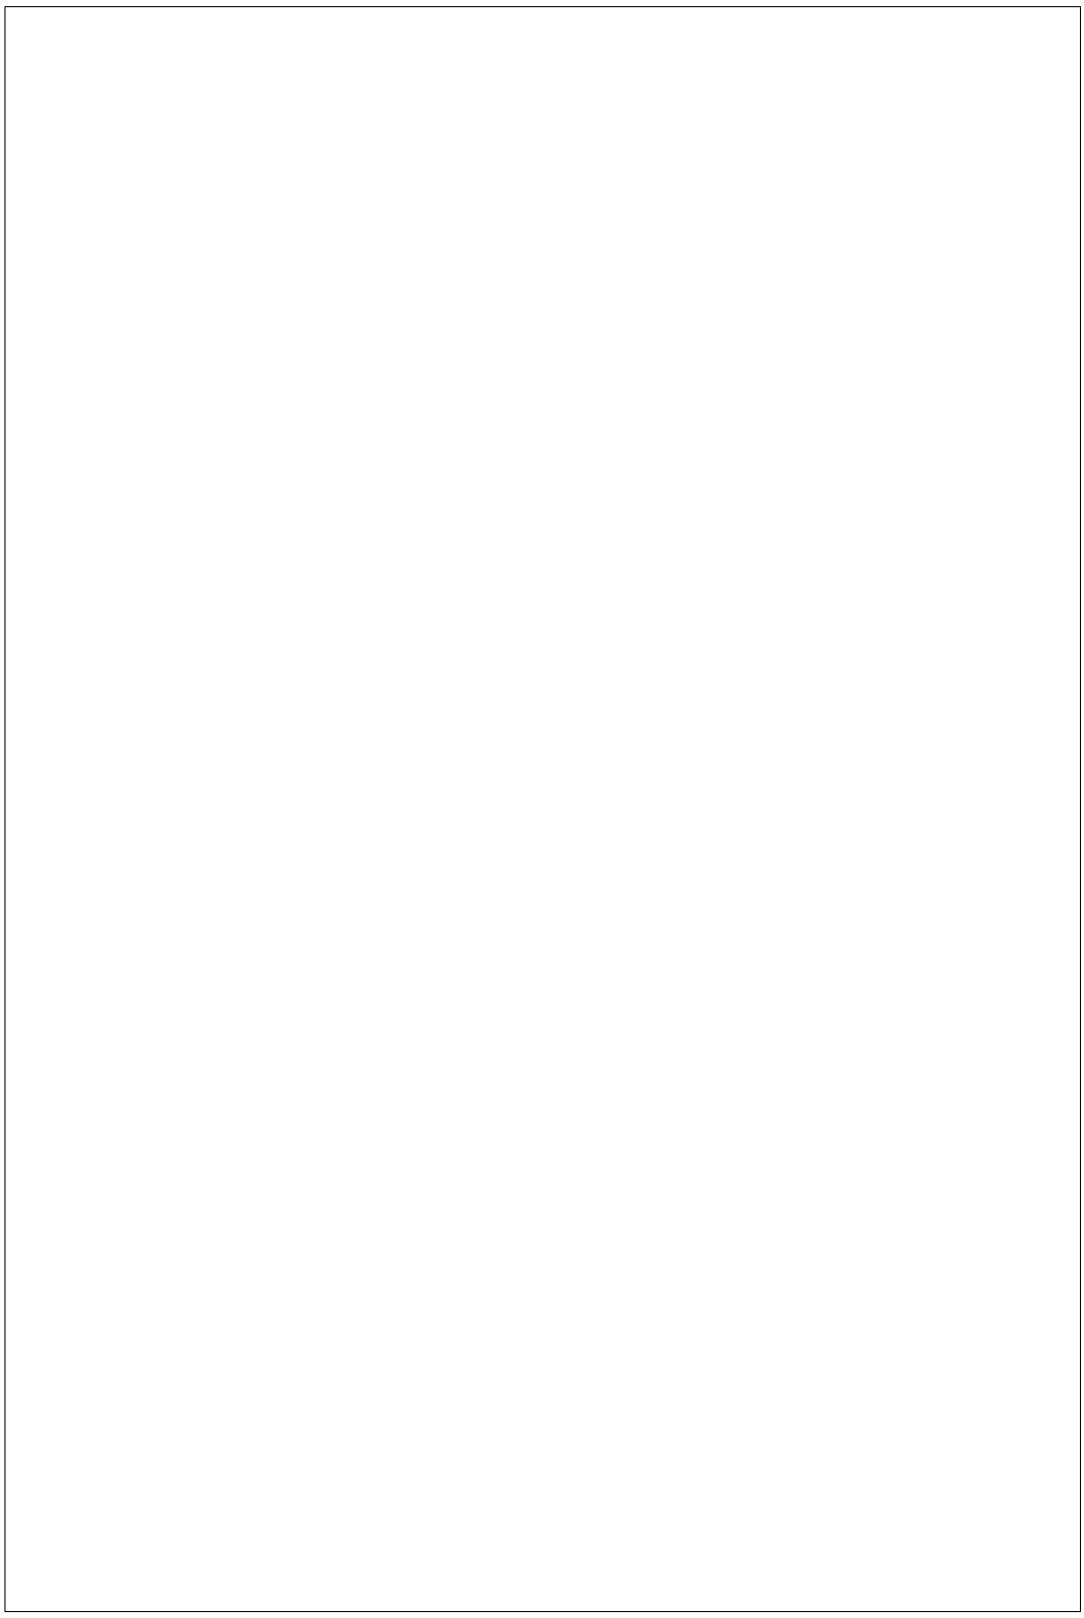

15.2 **Find the mass of the solid defined by the region contained within the cylinder $x^2 + y^2 = 1$ below the plane $z = 4$ and above the paraboloid $z = 1 - x^2 - y^2$. The density at any given point in the region is proportional to the distance from the axis of the cylinder.**

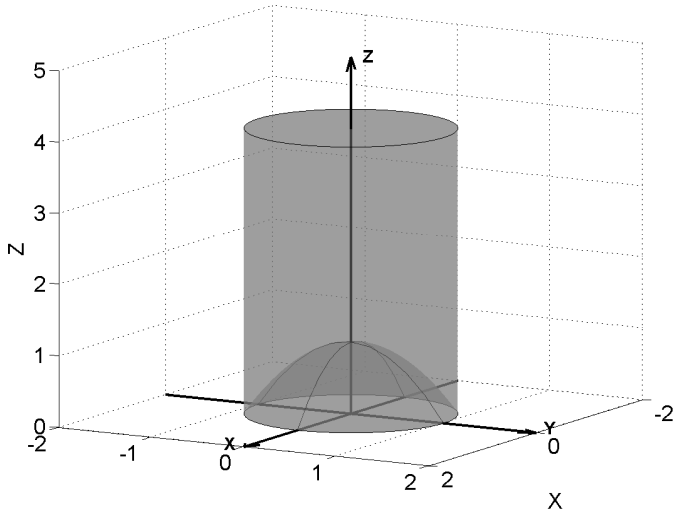

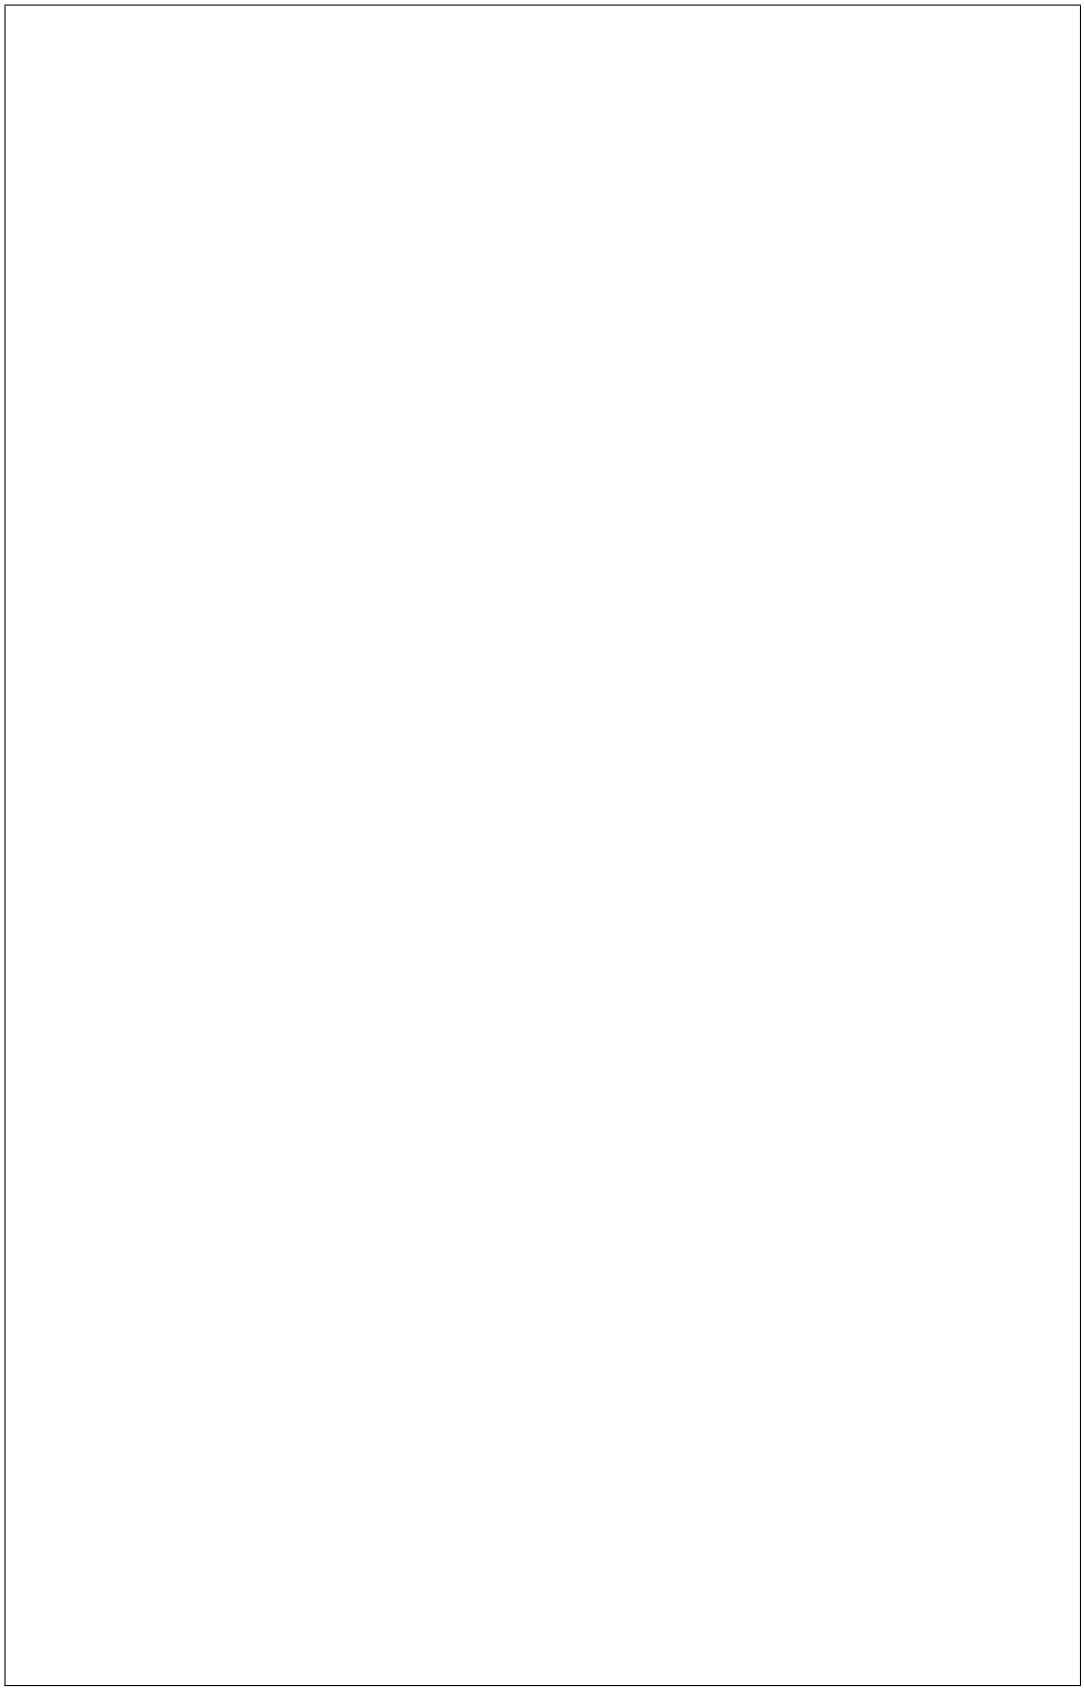

16 Spherical coordinates

By the end of this section, you should be able to answer the following questions:

- What is the relationship between rectangular coordinates and spherical coordinates?
- How do you transform a triple integral in rectangular coordinates into one in terms of spherical coordinates?
- What is the Jacobian of the transformation?

Sometimes it is useful to use spherical coordinates in order to simplify the integral. This involves the transformation

$$x = r\cos\theta\sin\phi, \quad y = r\sin\theta\sin\phi, \quad z = r\cos\phi. \tag{13}$$

In this case θ is longitude, ϕ is co-latitude, and r the distance from the origin.

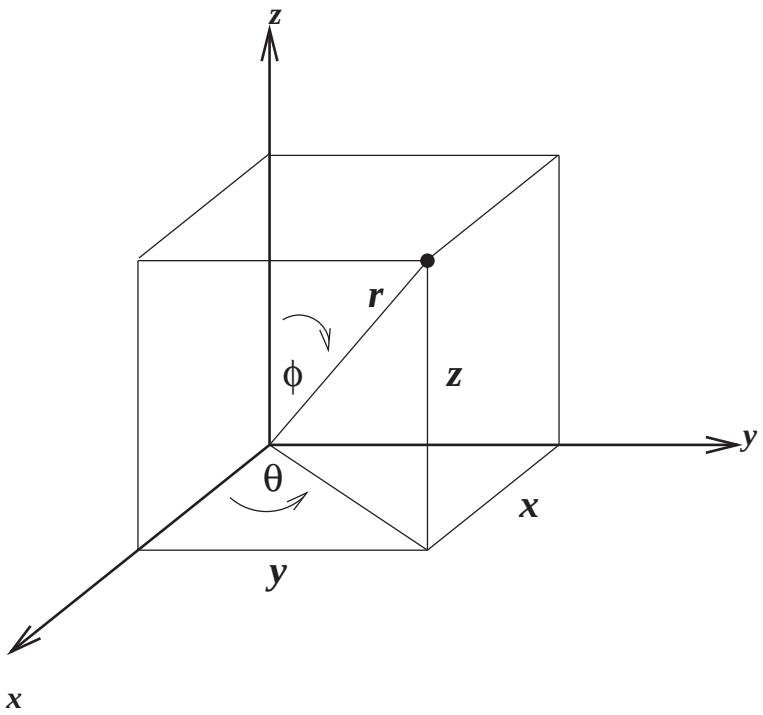

We now aim to calculate a small element of volume of a spherical shell. This will then show how in a triple integral we can transform from rectangular coordinates to spherical coordinates by substituting the transformation (13) and by making the change

$$dx\ dy\ dz \quad \longrightarrow \quad r^2 \sin\phi\ dr\ d\theta\ d\phi.$$

Consider the following diagram.

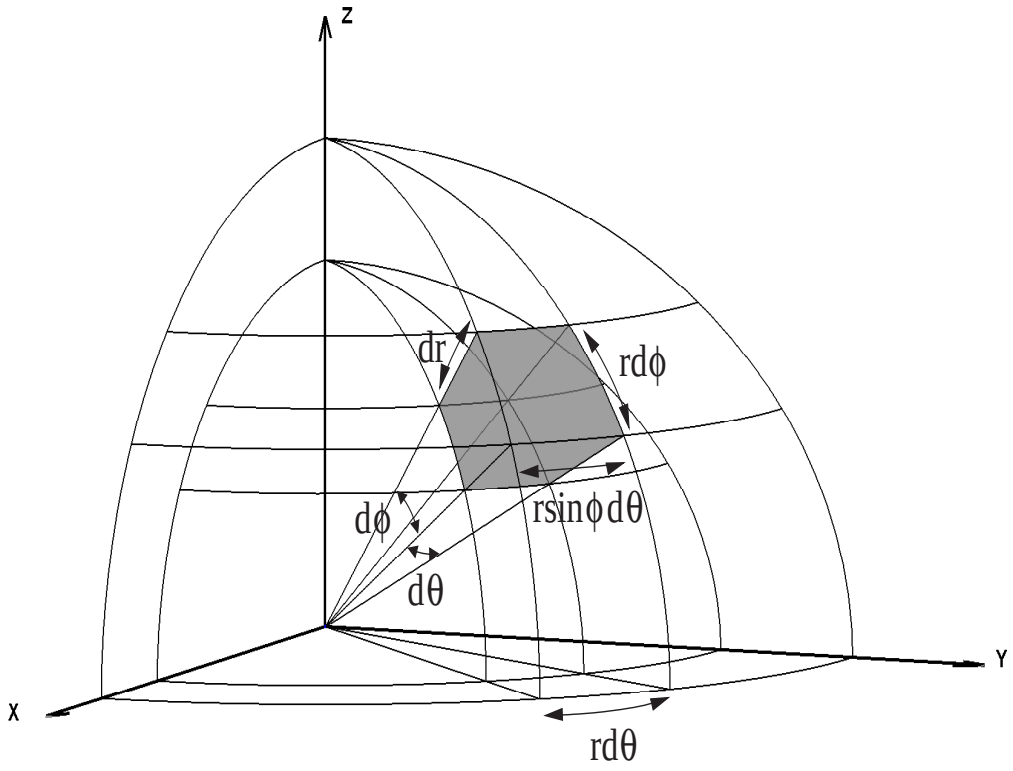

The important result is that the triple integral in rectangular coordinates transforms as follows:

$$\iiint_R f(x,y,z)\ dx\ dy\ dz$$
$$= \iiint_S f(r\cos\theta\sin\phi, r\sin\theta\sin\phi, r\cos\phi)\ r^2\sin\phi\ dr\ d\theta\ d\phi.$$

16.1 A simple example: Find the volume of a sphere of radius R.

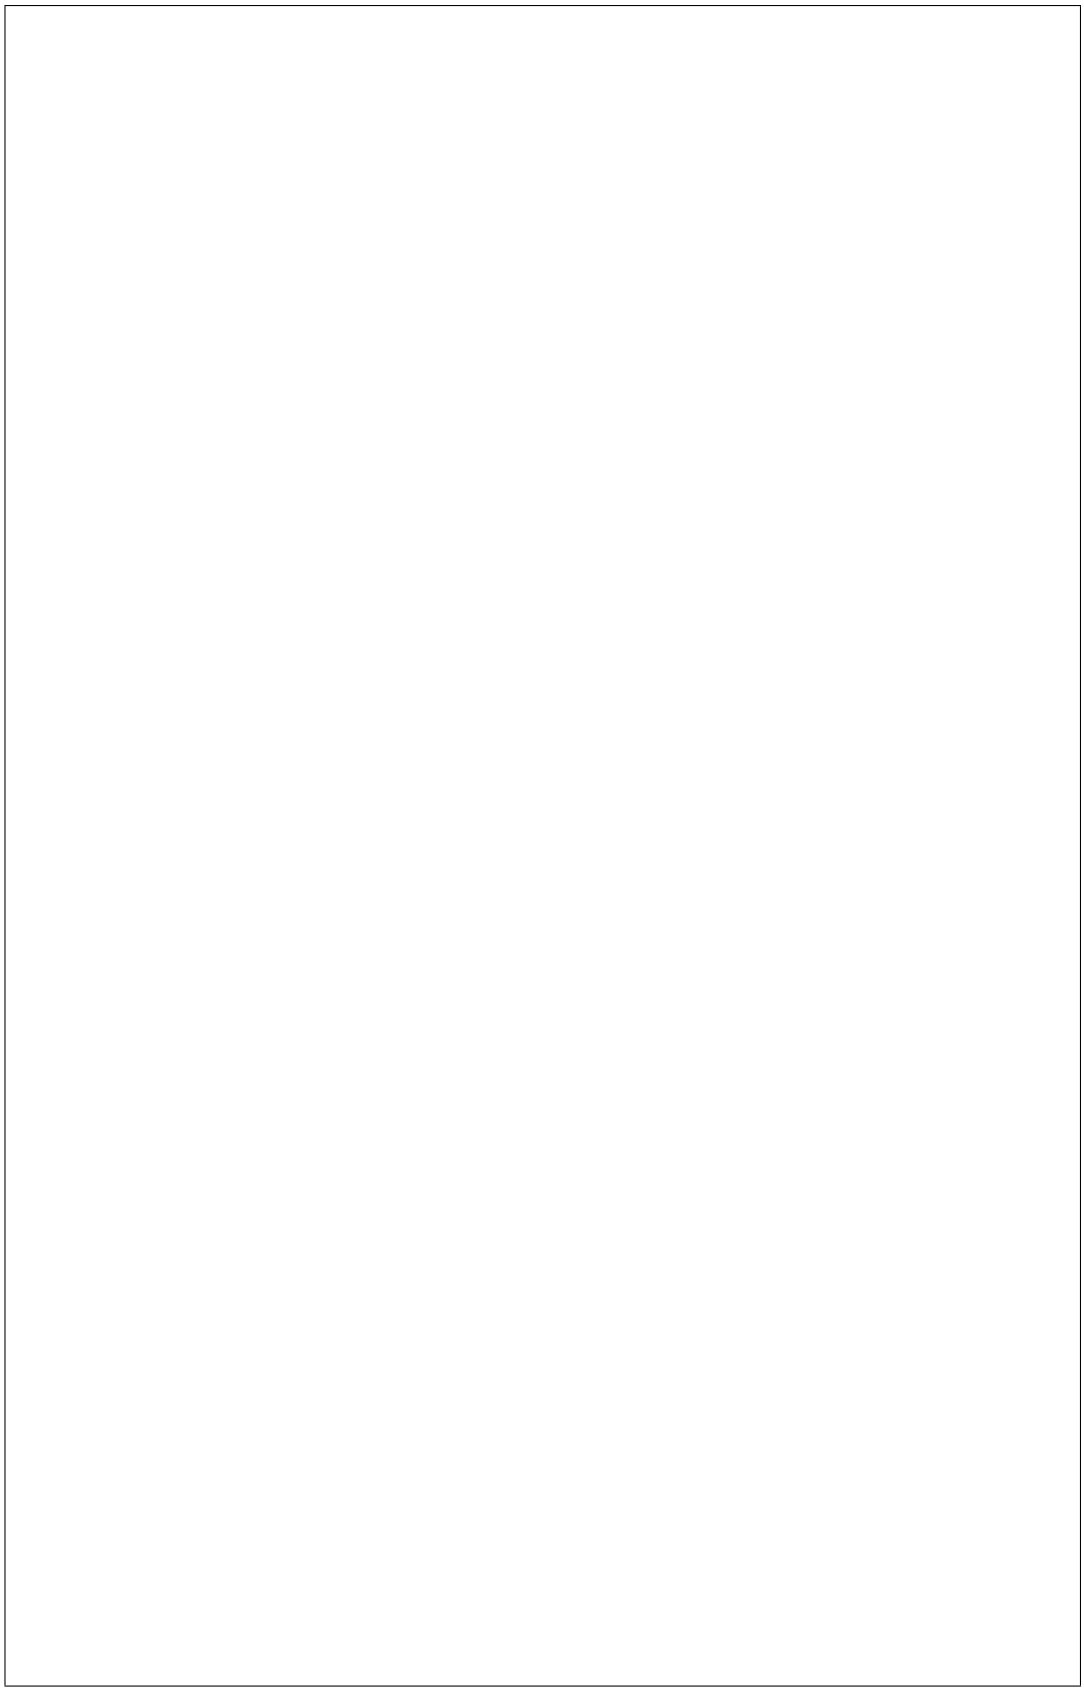

16.2 Find the mass of a sphere of radius R whose density is given by $\rho(x, y, z) = e^{-(x^2+y^2+z^2)^{1/2}}$.

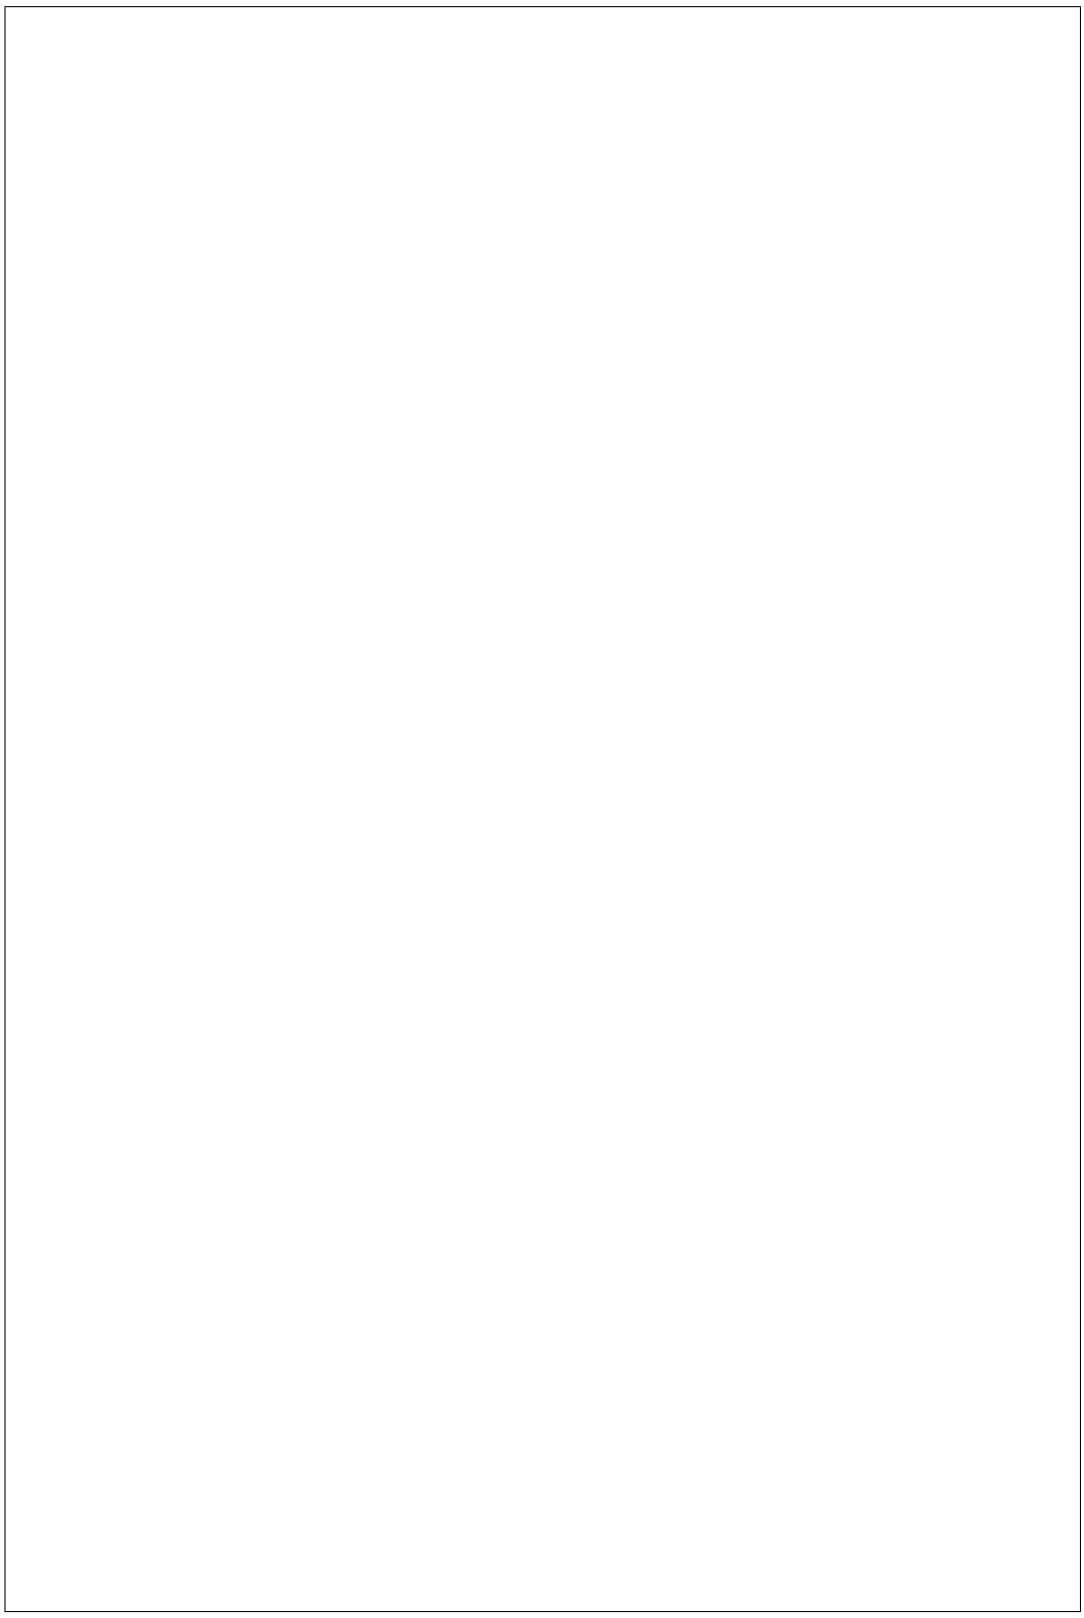

16.3 Find the volume of the "ice cream cone" R between a sphere of radius a (centred at the origin) and the cone $z = \sqrt{x^2 + y^2}$.

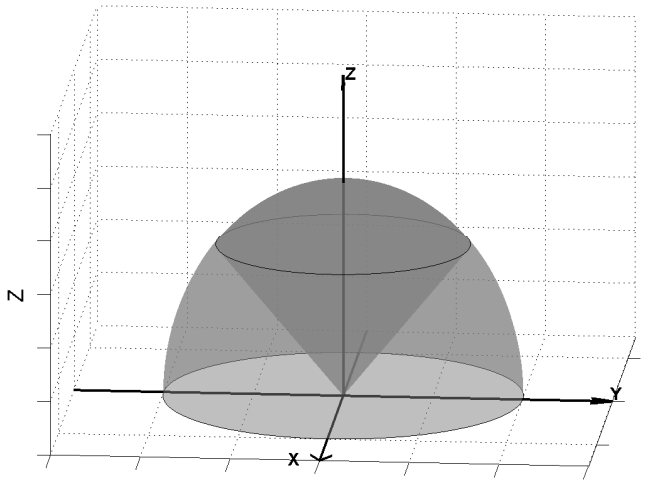

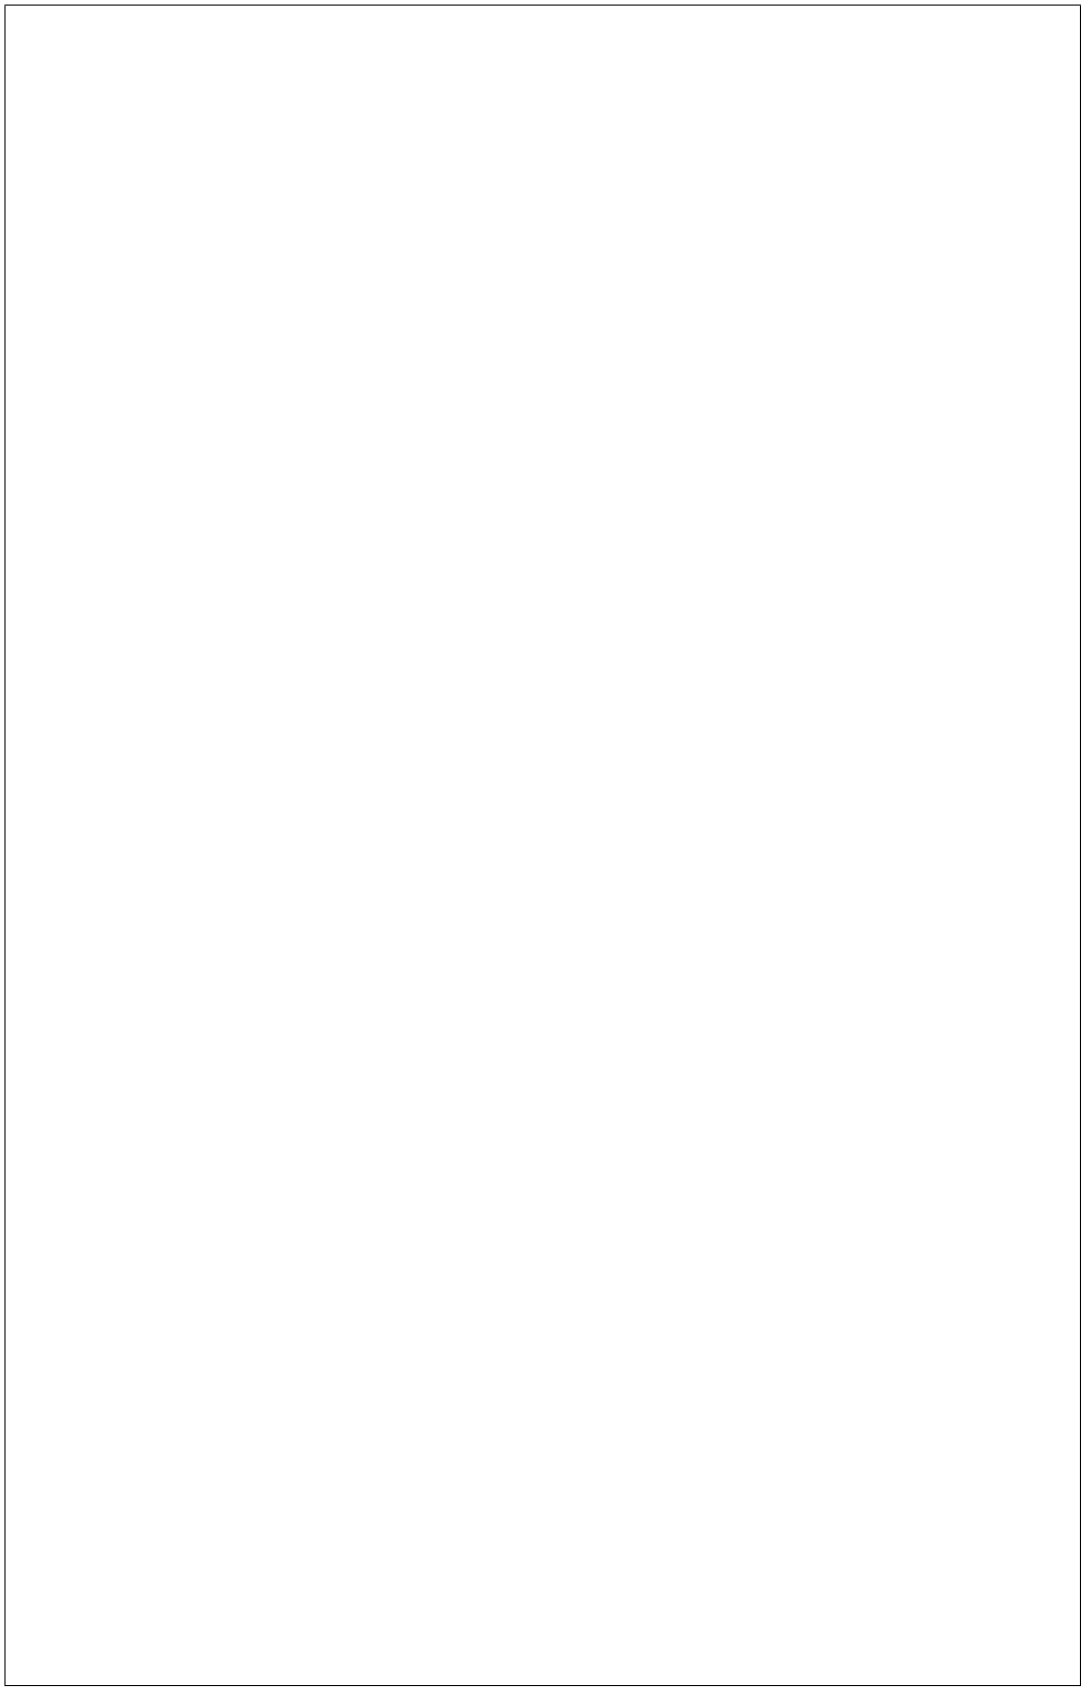

17 Moments of inertia (second moments)

By the end of this section, you should be able to answer the following questions:

- How do you locate the centre of mass of a solid object using a triple integral?
- How do you calculate the moments of inertia about the three coordinate axes?

The moment of inertia of a particle of mass m about an axis (x, y, or z) is defined to be mr^2 where r is the distance from the particle to the axis.

It is sometimes referred to as rotational inertia and can be thought of as the rotational analogue of mass for linear motion. For example, linear kinetic energy can be expressed as $\frac{1}{2}mv^2$, and the rotational kinetic energy as $\frac{1}{2}I\omega^2$. Linear momentum is determined by the formula $p = mv$, while angular momentum is given by $L = I\omega$. In these examples, I is the moment of inertia and ω the angular velocity.

As we have seen from previous examples, the mass of a solid with density $\rho(x, y, z)$ occupying a region R in $\mathbb{R}^3$ is given by

$$m = \iiint_R \rho(x, y, z)dV.$$

The moments about each of the three coordinate planes are

$$M_{yz} = \iiint_R x\rho(x, y, z)dV, \quad M_{xz} = \iiint_R y\rho(x, y, z)dV,$$

$$M_{xy} = \iiint_R z\rho(x, y, z)dV$$

The centre of mass is then located at the point $(\overline{x}, \overline{y}, \overline{z})$ where

$$\overline{x} = \frac{M_{yz}}{m}, \ \overline{y} = \frac{M_{xz}}{m}, \ \overline{z} = \frac{M_{xy}}{m}.$$

The moments of inertia about each of the three coordinate axes work out to be

$$I_x = \iiint_R (y^2 + z^2)\rho(x, y, z)dV,$$

$$I_y = \iiint_R (x^2 + z^2)\rho(x, y, z)dV,$$

$$I_z = \iiint_R (x^2 + y^2)\rho(x, y, z)dV.$$

17.1 Example: locate the centre of mass of a solid hemisphere of radius a with density proportional to the distance from the centre of the base. Find its moment of inertia about the z-axis.

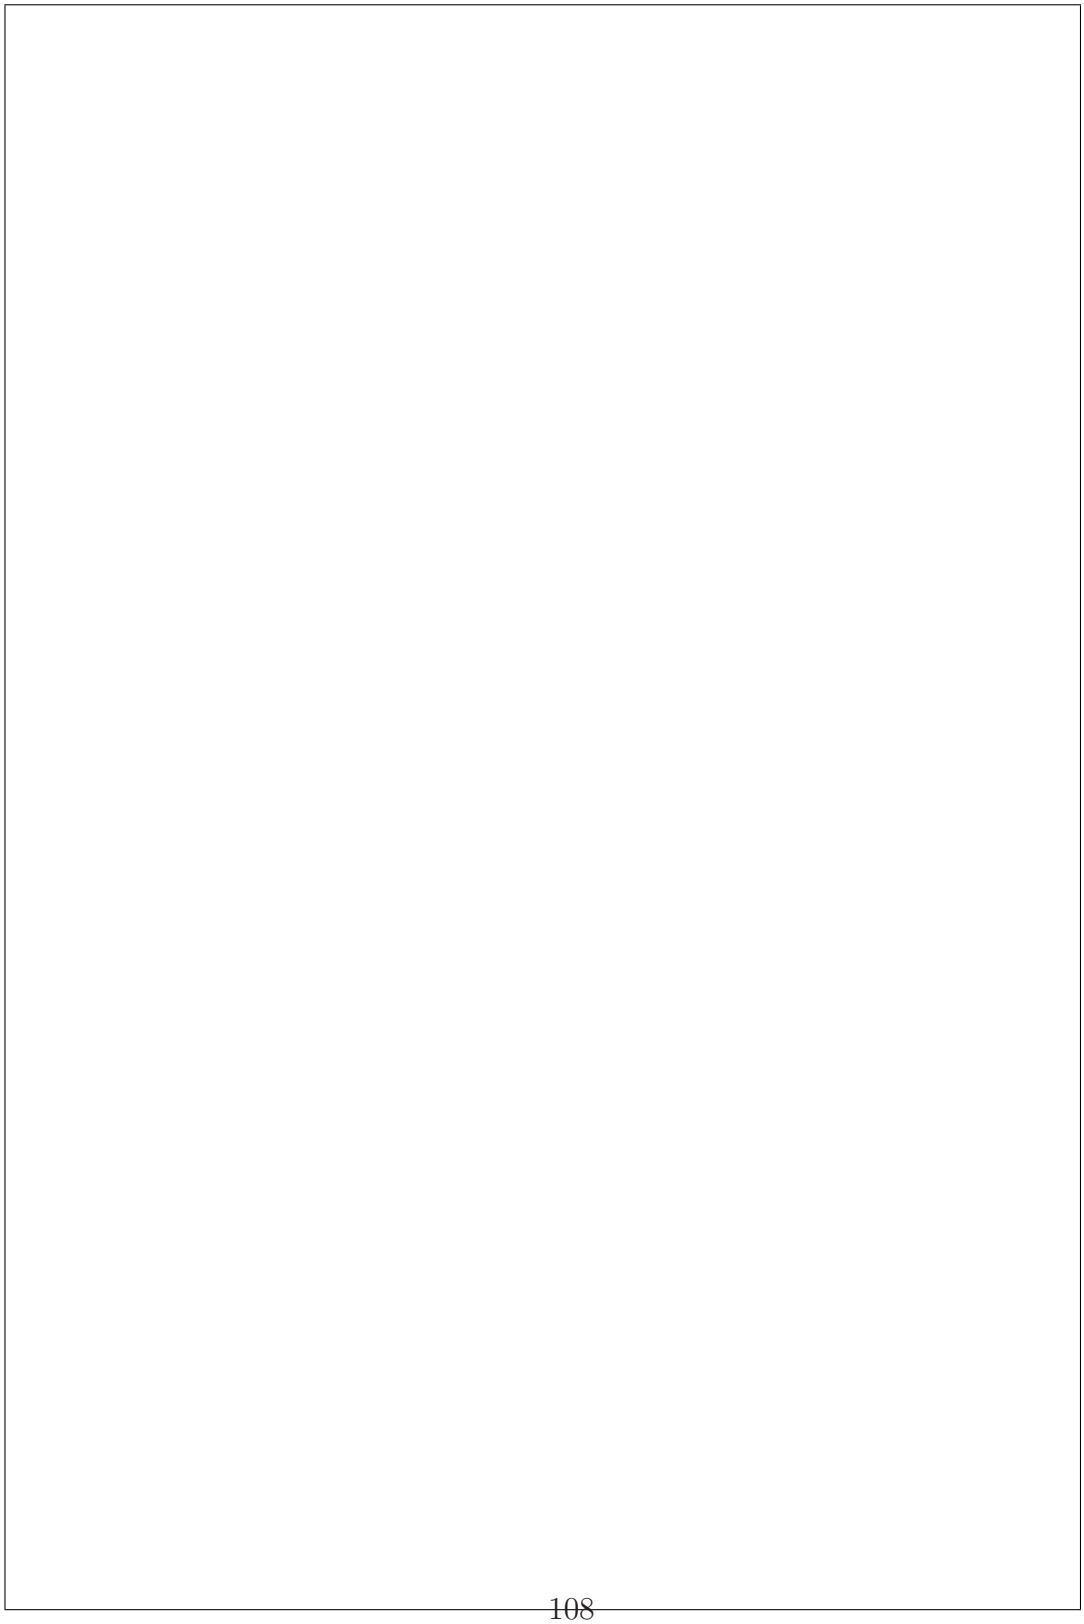

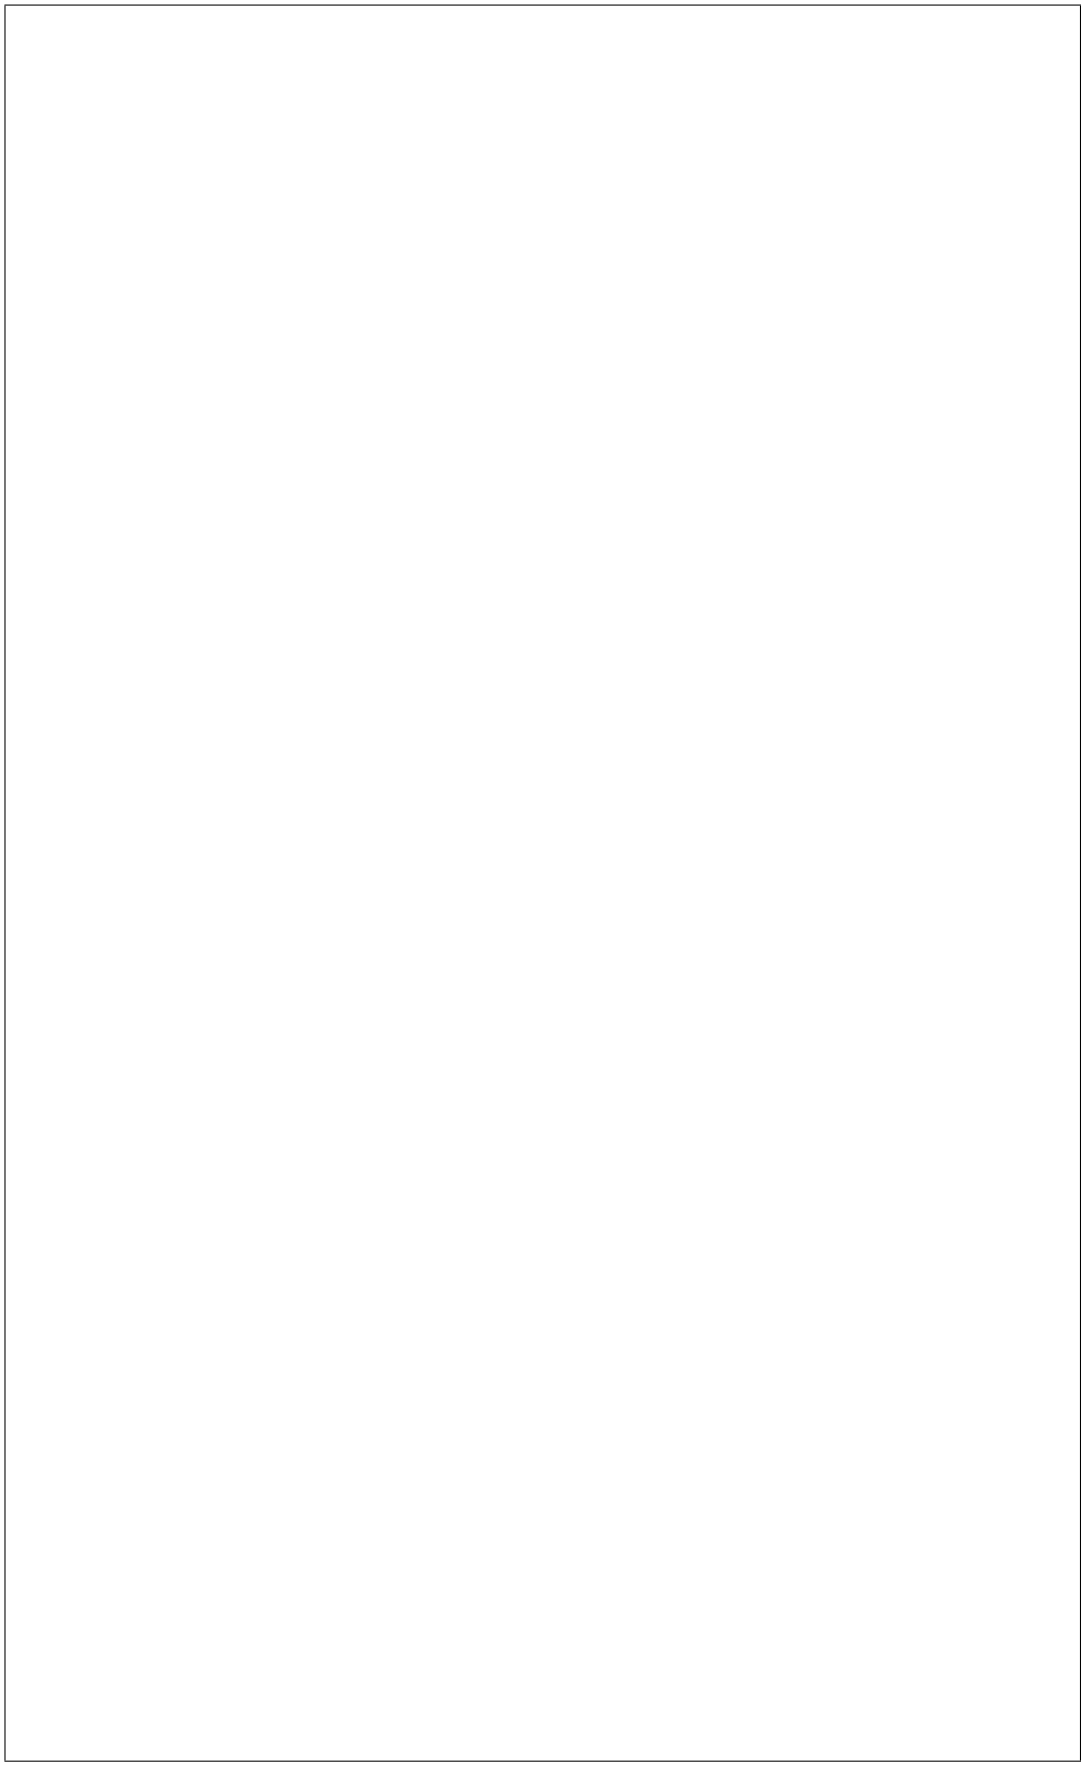

18 Conservative vector fields

By the end of this section, you should be able to answer the following questions:

- What is meant by a conservative vector field and a corresponding potential function?
- Given a potential function, how do you determine the corresponding conservative vector field?
- Given a conservative vector field, how do you determine a corresponding potential function?

18.1 Vector fields

In what follows, the notation is always

$$\boldsymbol{r} = x\boldsymbol{i} + y\boldsymbol{j} \text{ or } \boldsymbol{r} = x\boldsymbol{i} + y\boldsymbol{j} + z\boldsymbol{k}.$$

A vector field in the x-y plane is a vector function of 2 variables

$$\begin{aligned} \boldsymbol{F}(\boldsymbol{r}) = \boldsymbol{F}(x,y) &= (F_1(x,y), F_2(x,y)) \\ &= F_1(x,y)\boldsymbol{i} + F_2(x,y)\boldsymbol{j}. \end{aligned}$$

That is, associated to a point (x, y) is the vector $\boldsymbol{F}(\boldsymbol{r})$.

18.1.1 Example: $\boldsymbol{F}(\boldsymbol{r}) = (-y, x) = -y\boldsymbol{i} + x\boldsymbol{j}$.

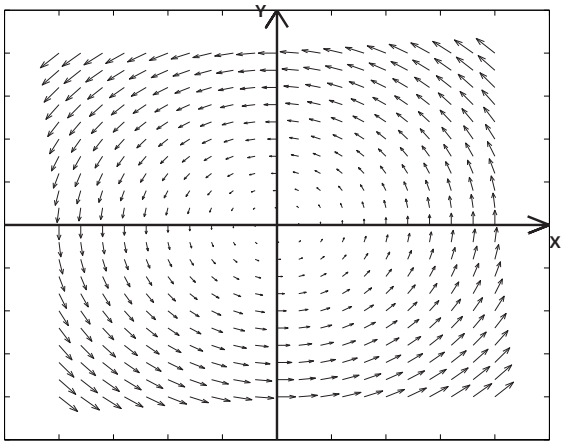

Similarly a vector field in 3-D is a vector function of 3 variables

$$\begin{aligned}\boldsymbol{F}(\boldsymbol{r}) &= \boldsymbol{F}(x,y,z)\\ &= (F_1(x,y,z), F_2(x,y,z), F_3(x,y,z))\\ &= F_1(x,y,z)\boldsymbol{i} + F_2(x,y,z)\boldsymbol{j} + F_3(x,y,z)\boldsymbol{k}\end{aligned}$$

18.1.2 Example: Newtonian gravitational field

$$\begin{aligned}\boldsymbol{F}(\boldsymbol{r}) &= -\frac{mMG}{|\boldsymbol{r}|^3}\boldsymbol{r} = \boldsymbol{F}(x,y,z)\\ &= \frac{-mMGx}{(x^2+y^2+z^2)^{3/2}}\boldsymbol{i} + \frac{-mMGy}{(x^2+y^2+z^2)^{3/2}}\boldsymbol{j}\\ &\quad + \frac{-mMGz}{(x^2+y^2+z^2)^{3/2}}\boldsymbol{k}\end{aligned}$$

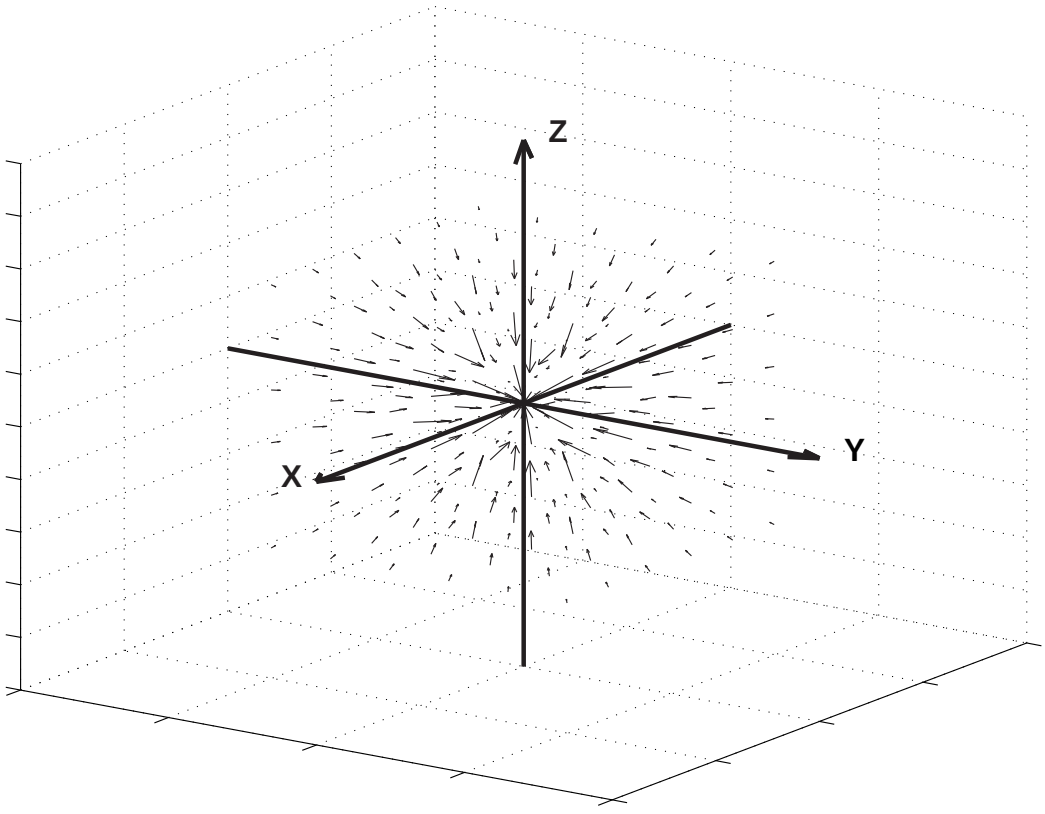

18.2 Gradient of a scalar field, conservative vector fields

Recall for a differentiable scalar function $f(x, y)$ in two dimensions, we define

$$\text{grad} f = \frac{\partial f}{\partial x}\boldsymbol{i} + \frac{\partial f}{\partial y}\boldsymbol{j}.$$

For a differentiable scalar function $f(x, y, z)$ in three dimensions, we define

$$\text{grad} f = \frac{\partial f}{\partial x}\boldsymbol{i} + \frac{\partial f}{\partial y}\boldsymbol{j} + \frac{\partial f}{\partial z}\boldsymbol{k}.$$

Alternatively we define the differential operator

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

so $\text{grad} f = \nabla f$.

18.2.1 Example: find the gradient of $f(x, y, z) = x^2y^3z^4$.

Note ∇f is a vector. It's length and direction are independent of the choice of coordinates. ∇f (evaluated at a given point P) is in the direction of maximum increase of f at P.

You may see the scalar function f referred to as a scalar field. If a vector field $\boldsymbol{v}$ and a scalar field f are related by $\boldsymbol{v} = \nabla f$, we call f a potential function and $\boldsymbol{v}$ a *conservative vector field.*

18.2.2 Verify that the Newtonian gravitational field is conservative with potential function $f(x, y, z) = \dfrac{mMG}{\sqrt{x^2 + y^2 + z^2}}$.

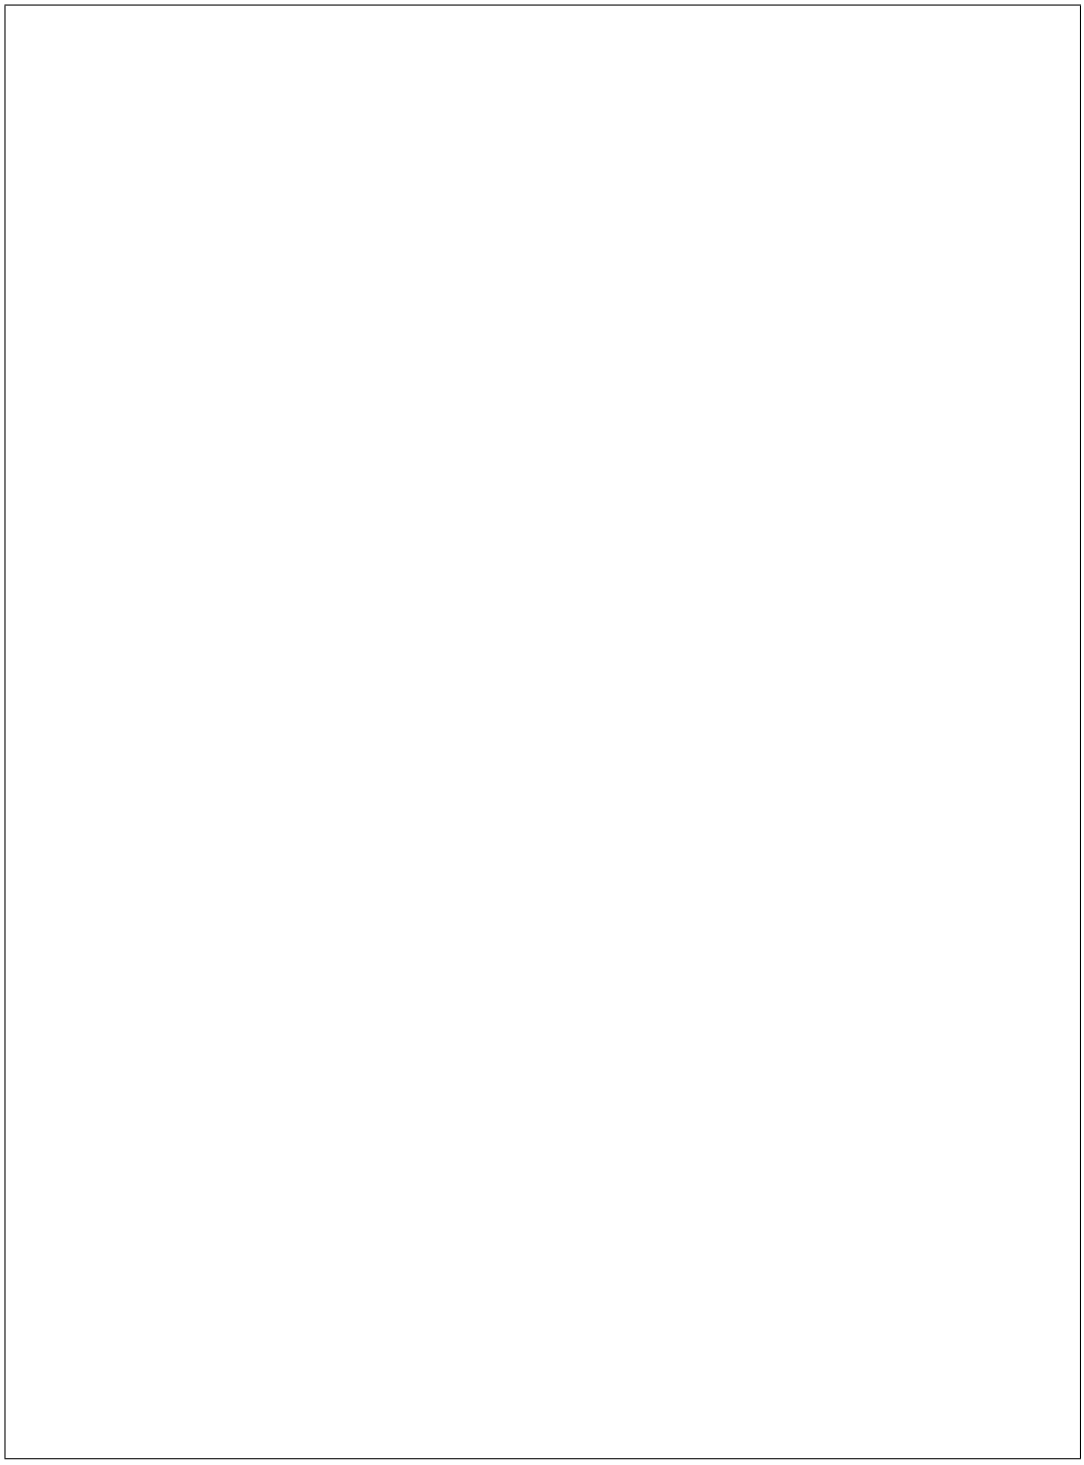

Given a conservative vector field, how can we determine a corresponding potential function? The next example outlines this procedure.

18.2.3 The vector field $\boldsymbol{F}(x, y) = (3 + 2xy)\boldsymbol{i} + (x^2 - 3y^2)\boldsymbol{j}$ is conservative. Find a corresponding potential function.

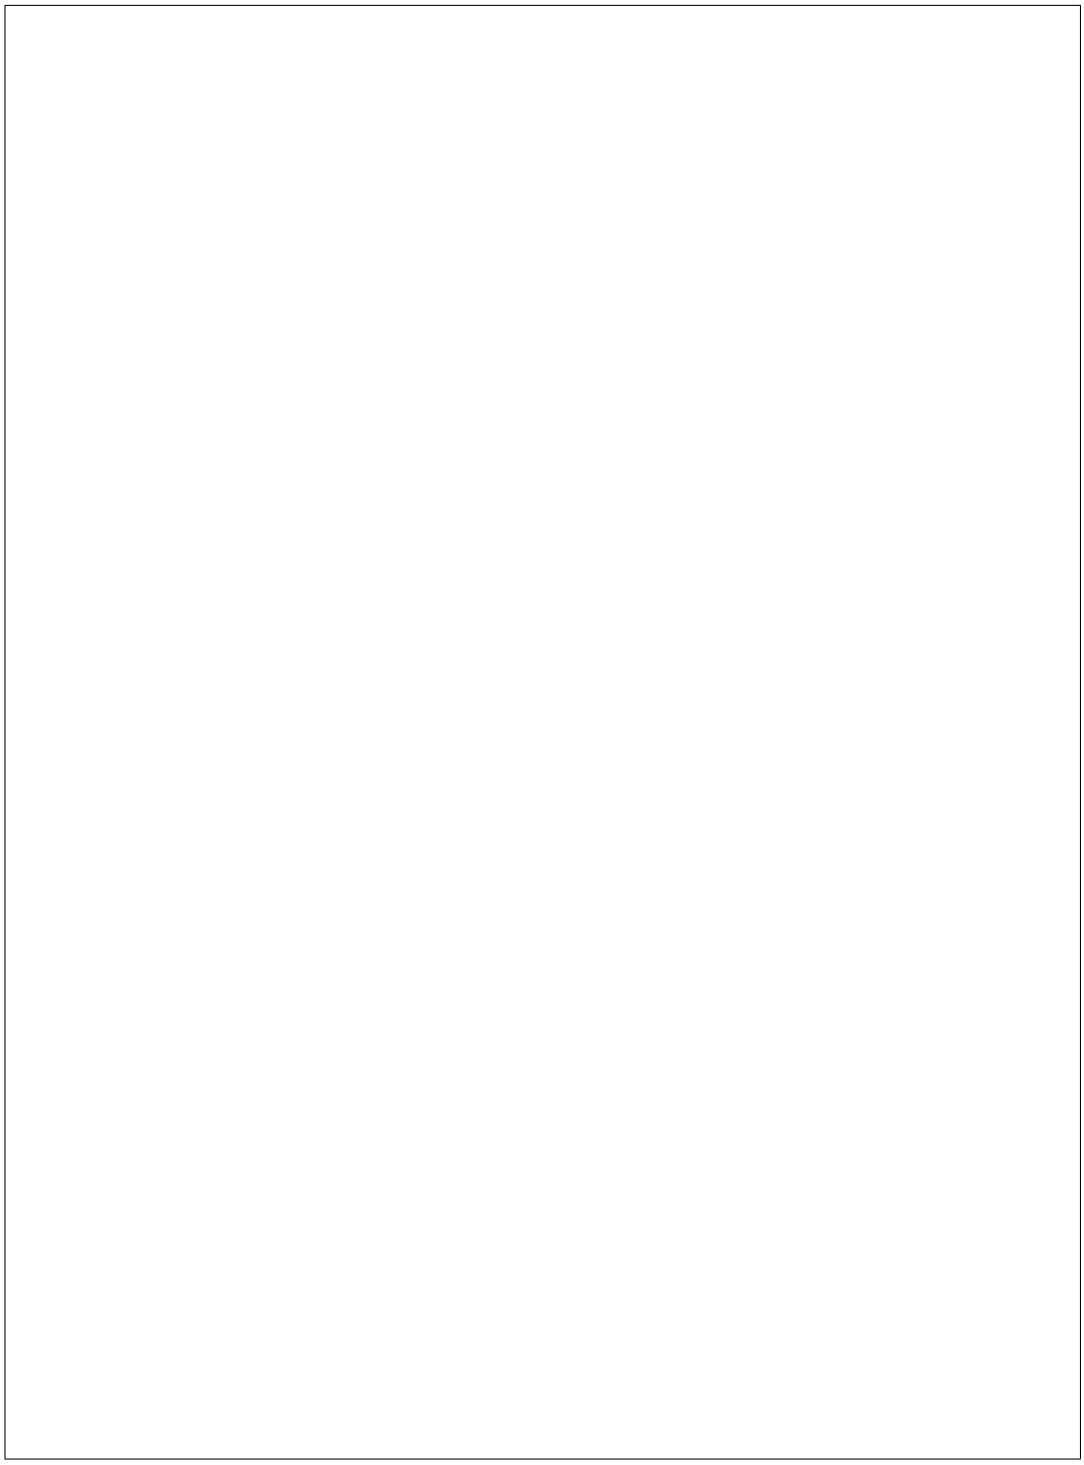

Can we still determine a potential function when the conservative vector field is in three dimensions?

18.2.4 The vector field $\boldsymbol{F}(x, y, z) = y^2\boldsymbol{i}+(2xy+e^{3z})\boldsymbol{j}+3ye^{3z}\boldsymbol{k}$ is conservative. Find a corresponding potential function.

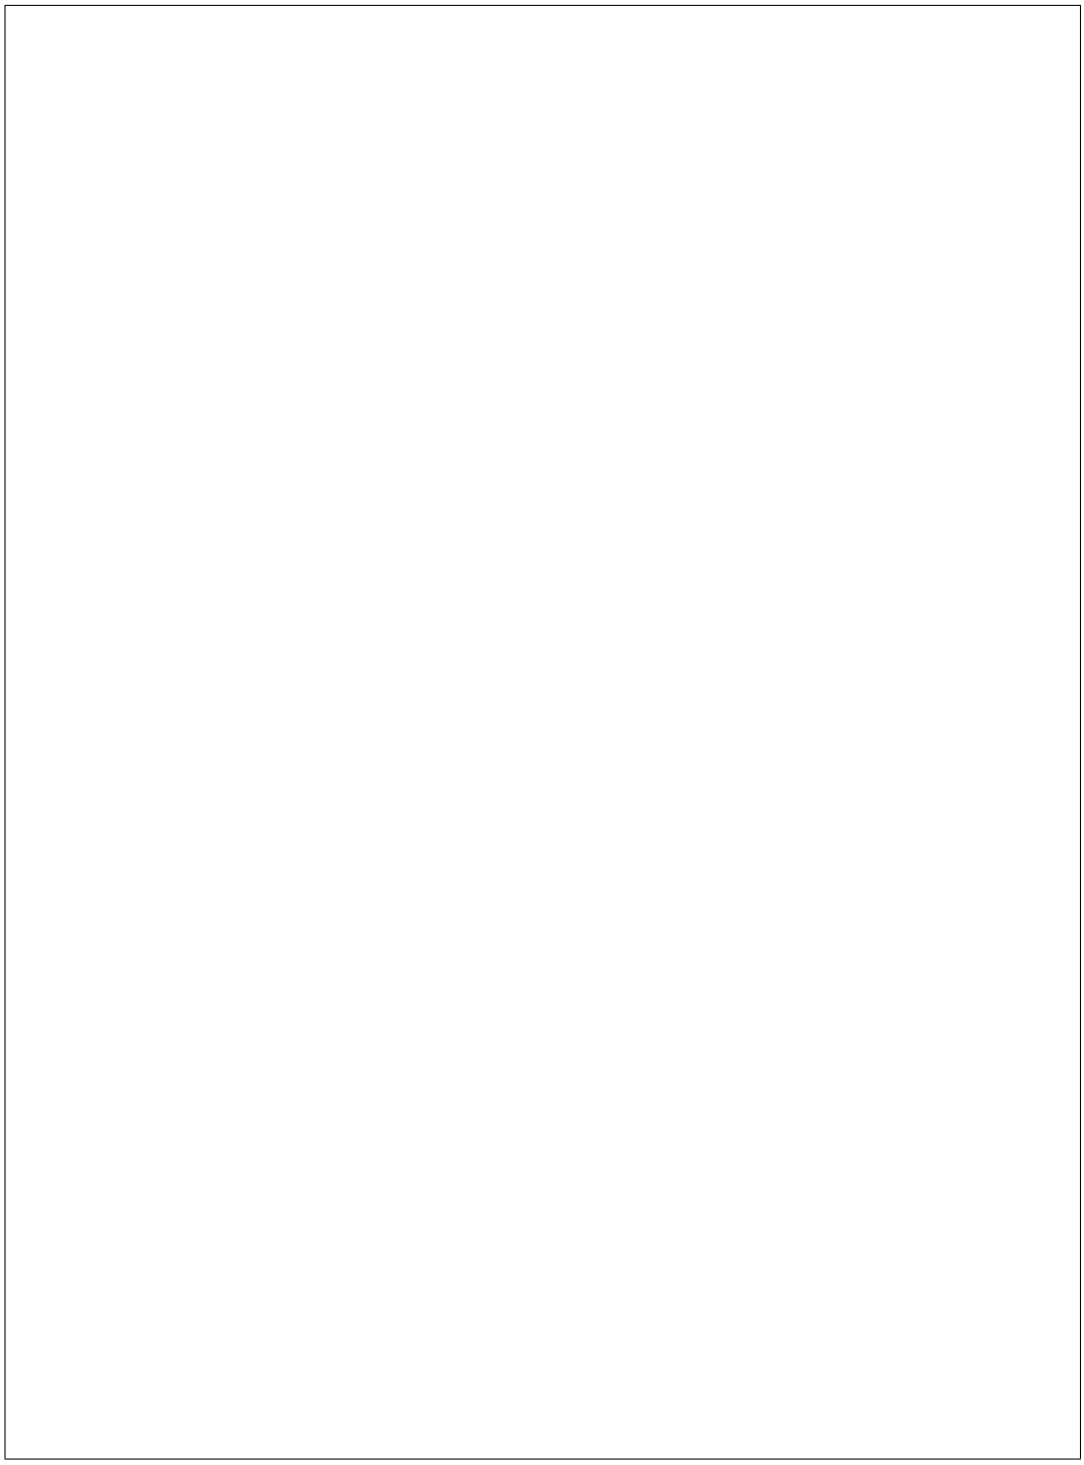

Is there a way of determining whether or not a given vector field is conservative? To answer this question, we need to go back to the study of line integrals.

19 The fundamental theorem for line integrals, path independence

By the end of this section, you should be able to answer the following questions:

- How do you evaluate line integrals?
- What is the fundamental theorem for line integrals and its consequences?
- What is a path independent line integral and what are its connections with conservative vector fields and line integrals over closed curves?

19.1 Line integrals in the plane

Recall the definite integral $\int_a^b f(x)\ dx$ gives the net area above the x-axis and below its image $y = f(x)$. We can generalise this.

Consider the following problem: How do we calculate the area of the region between the curve C in the x-y plane and its image on the surface $z = f(x, y)$?

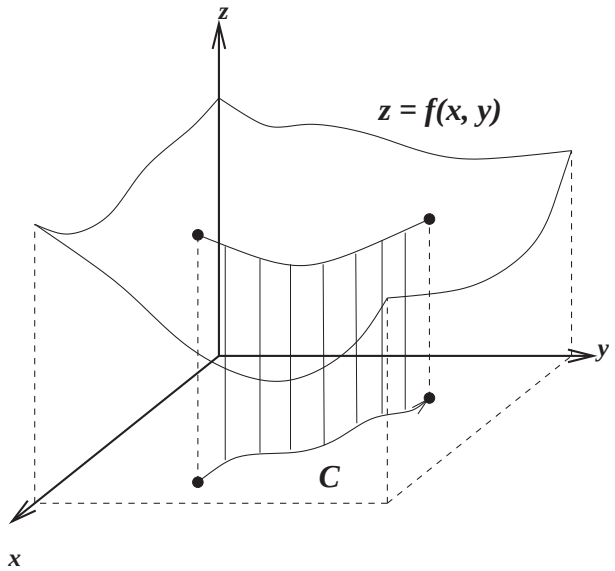

If the curve C can be parametrised by $\boldsymbol{r}(t) = x(t)\boldsymbol{i} + y(t)\boldsymbol{j}$ for $a \leq t \leq b$, then the area is given by the formula

$$\text{area } = \int_C f(x, y)\ dS = \int_a^b f(x(t), y(t))|\boldsymbol{r}'(t)|\ dt,$$

where dS is the infinitesimal element of arclength of C.

19.2 Line integrals of vector fields

We can also consider integrating a vector field over a curve in the plane.

Let C be a piecewise continuous smooth curve in the x-y plane joining points A and B. Let $\boldsymbol{F}(x,y) = F_1(x,y)\boldsymbol{i} + F_2(x,y)\boldsymbol{j}$ be a vector field. A line integral is given by

$$\begin{aligned}\int_C \boldsymbol{F}(\boldsymbol{r}) \cdot \boldsymbol{dr} &= \int_C (F_1(x,y)dx + F_2(x,y)dy) \\ &= \int_a^b \boldsymbol{F}(\boldsymbol{r}(t)) \cdot \boldsymbol{r}'(t)\ dt\end{aligned}$$

where $\boldsymbol{r} = x\boldsymbol{i} + y\boldsymbol{j}$, $\boldsymbol{dr} = dx\boldsymbol{i} + dy\boldsymbol{j}$ and x, y are parameterised by $t \in [a,b]$.

Note that we can also write the line integral as $\displaystyle\int_C \boldsymbol{F}(x,y) \cdot \boldsymbol{T}(x,y)\ dS$ where $\boldsymbol{T}$ is a unit tangent vector to the curve C at the point (x,y) on C.

In the case $\boldsymbol{F}$ is a field of force, you should already be able to determine the work done by $\boldsymbol{F}$ in moving a particle along a curve C. Namely, you should already know that

$$\begin{aligned}\text{work} &= \int_a^b \left[\boldsymbol{F}(\boldsymbol{r}(t)) \cdot \frac{\boldsymbol{r}'(t)}{|\boldsymbol{r}'(t)|}\right] |\boldsymbol{r}'(t)|dt \\ &= \int_a^b \boldsymbol{F}(\boldsymbol{r}(t)) \cdot \boldsymbol{r}'(t)dt.\end{aligned}$$

19.3 Evaluating line integrals

In general, to evaluate a line integral

$$\int_C f(x,y)\ dS,$$

which includes line integrals of the form

$$\int_C \boldsymbol{F} \cdot \boldsymbol{dr} = \int_C \boldsymbol{F} \cdot \frac{\boldsymbol{r}'}{|\boldsymbol{r}'|}\ dS,$$

we start by parametrising C with $\boldsymbol{r}(t)$ and in the integral replace dS by $|\boldsymbol{r}'(t)|\ dt$. Then evaluate the integral as a definite integral in t. The bounds of integration for t are those values corresponding to the endpoints of C.

19.3.1 Example: let $A = (0, 1)$, $B = (1, 2)$. Evaluate $\int_C ((x^2 - y)dx + (y^2 + x)dy)$ along the curve C given by: (i) the straight line from A to B; (ii) the parabola $y = x^2 + 1$ from A to B.

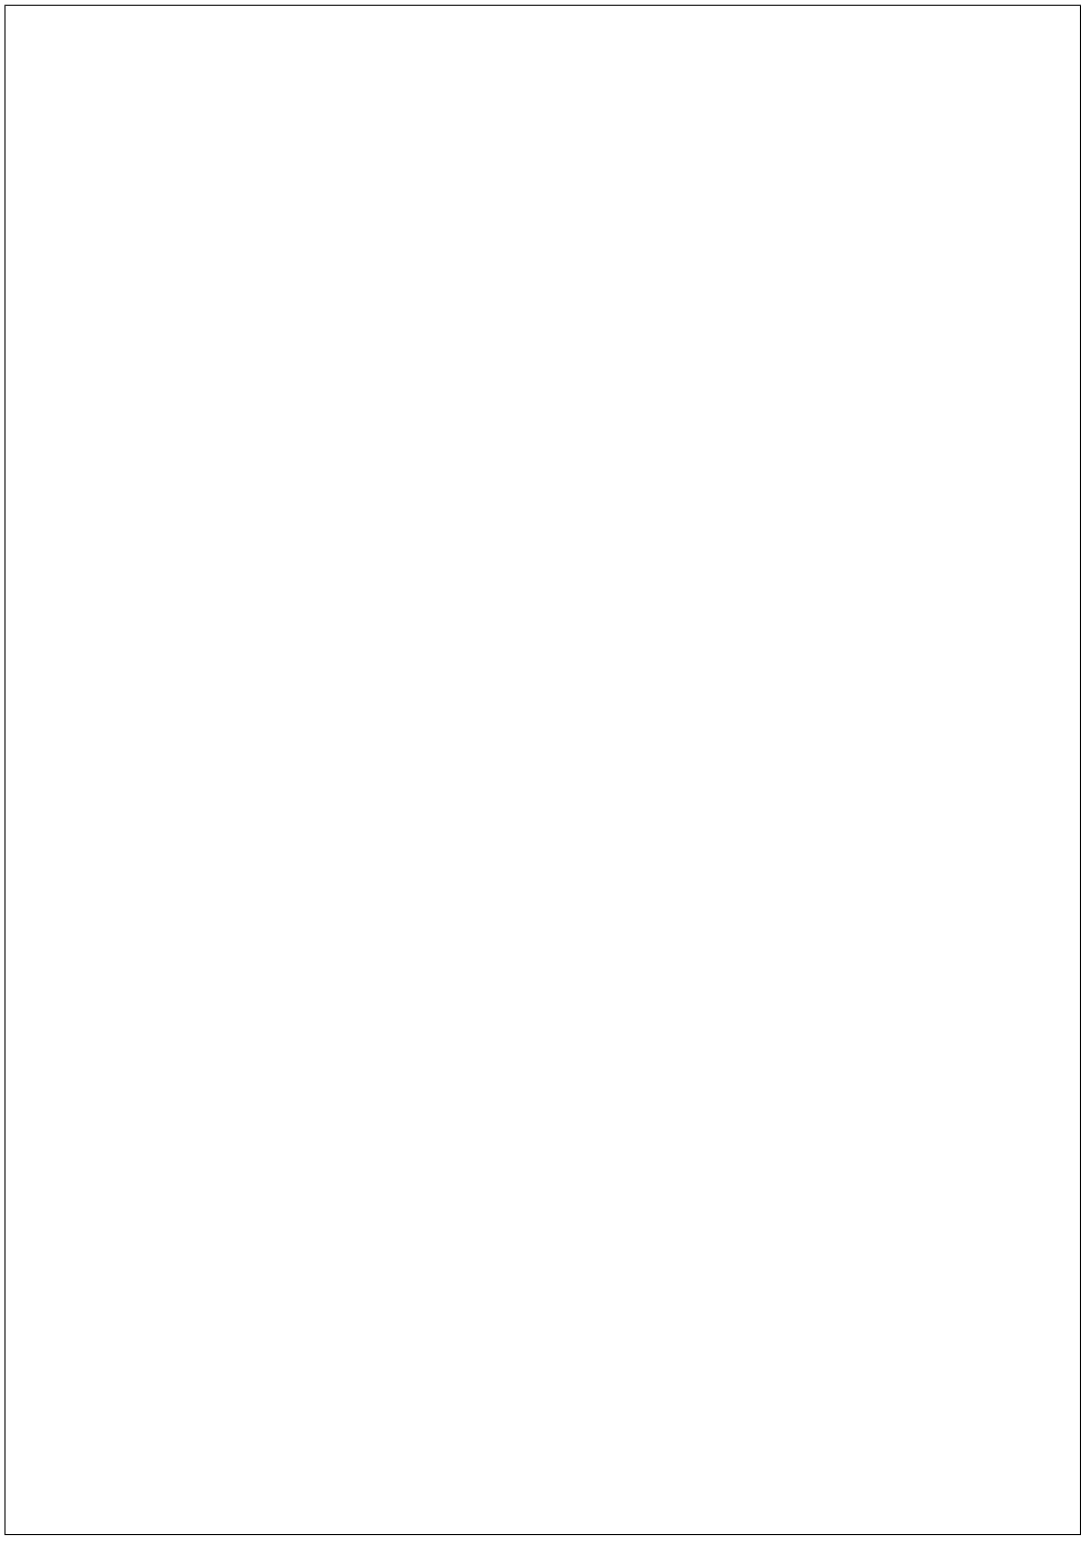

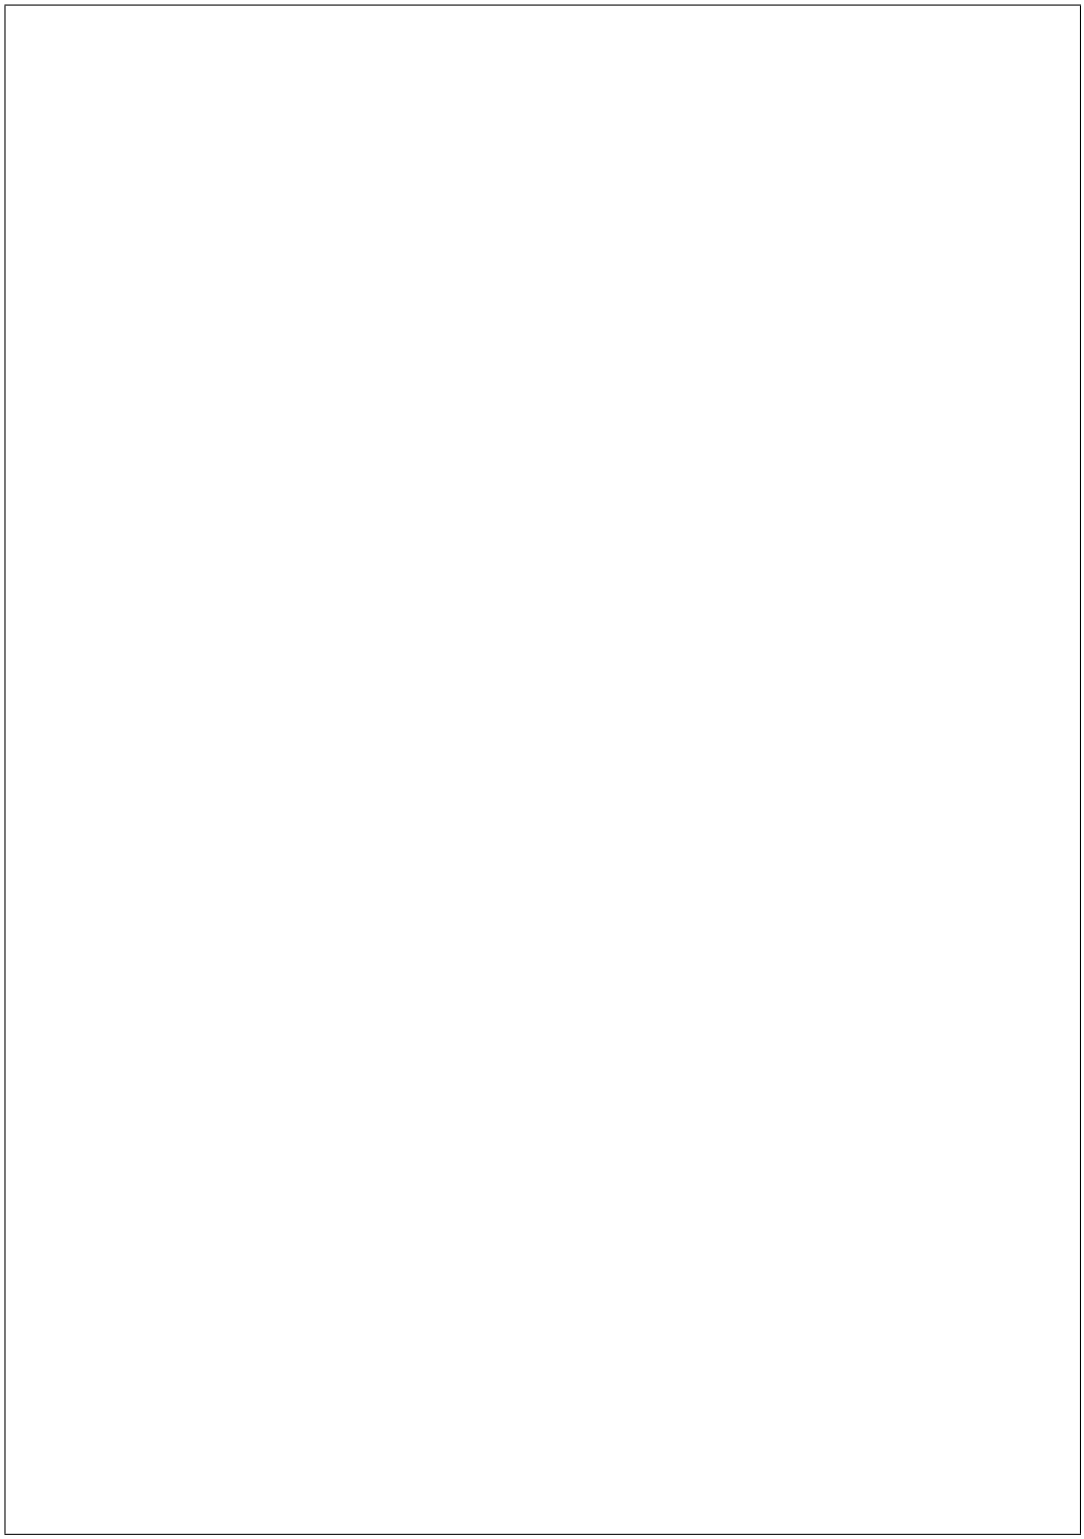

Note the line integrals in the previous example were *path dependent*. In other words, they have different values for different paths.

We will now investigate path independent line integrals.

19.4 Line integrals of conservative vector fields, path independence.

If $\boldsymbol{F}$ is a continuous vector field with domain D, we say the line integral $\int_C \boldsymbol{F} \cdot \boldsymbol{dr}$ is *path independent* if

$$\int_{C_1} \boldsymbol{F} \cdot \boldsymbol{dr} = \int_{C_2} \boldsymbol{F} \cdot \boldsymbol{dr}$$

for any two paths C_1 and C_2 in D that have the same end points.

19.4.1 The fundamental theorem for line integrals

If C is a smooth curve determined by $\boldsymbol{r}(t)$ for $t \in [a, b]$ and $f(x, y)$ is differentiable with ∇f being continuous on C, then

$$\int_C \nabla f \cdot \boldsymbol{dr} = f(\boldsymbol{r}(b)) - f(\boldsymbol{r}(a)).$$

Proof:

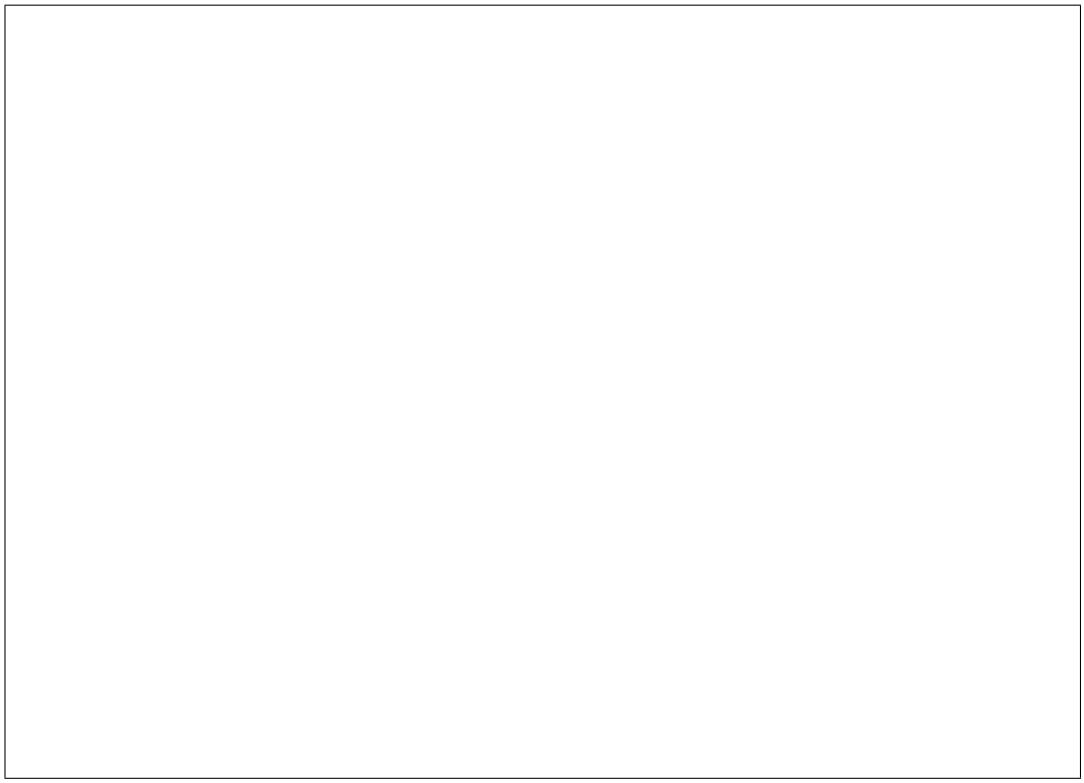

One consequence is that for conservative vector fields ∇f, we have

$$\int_{C_1} \nabla f \cdot \boldsymbol{dr} = \int_{C_2} \nabla f \cdot \boldsymbol{dr}.$$

That is, the line integral of a conservative vector field is path independent.

It turns out, the converse is also true. Suppose $\boldsymbol{F}$ is continuous on an open, connected region D. If $\displaystyle\int_C \boldsymbol{F} \cdot \boldsymbol{dr}$ is path independent in D, then $\boldsymbol{F}$ is conservative.

Proof:

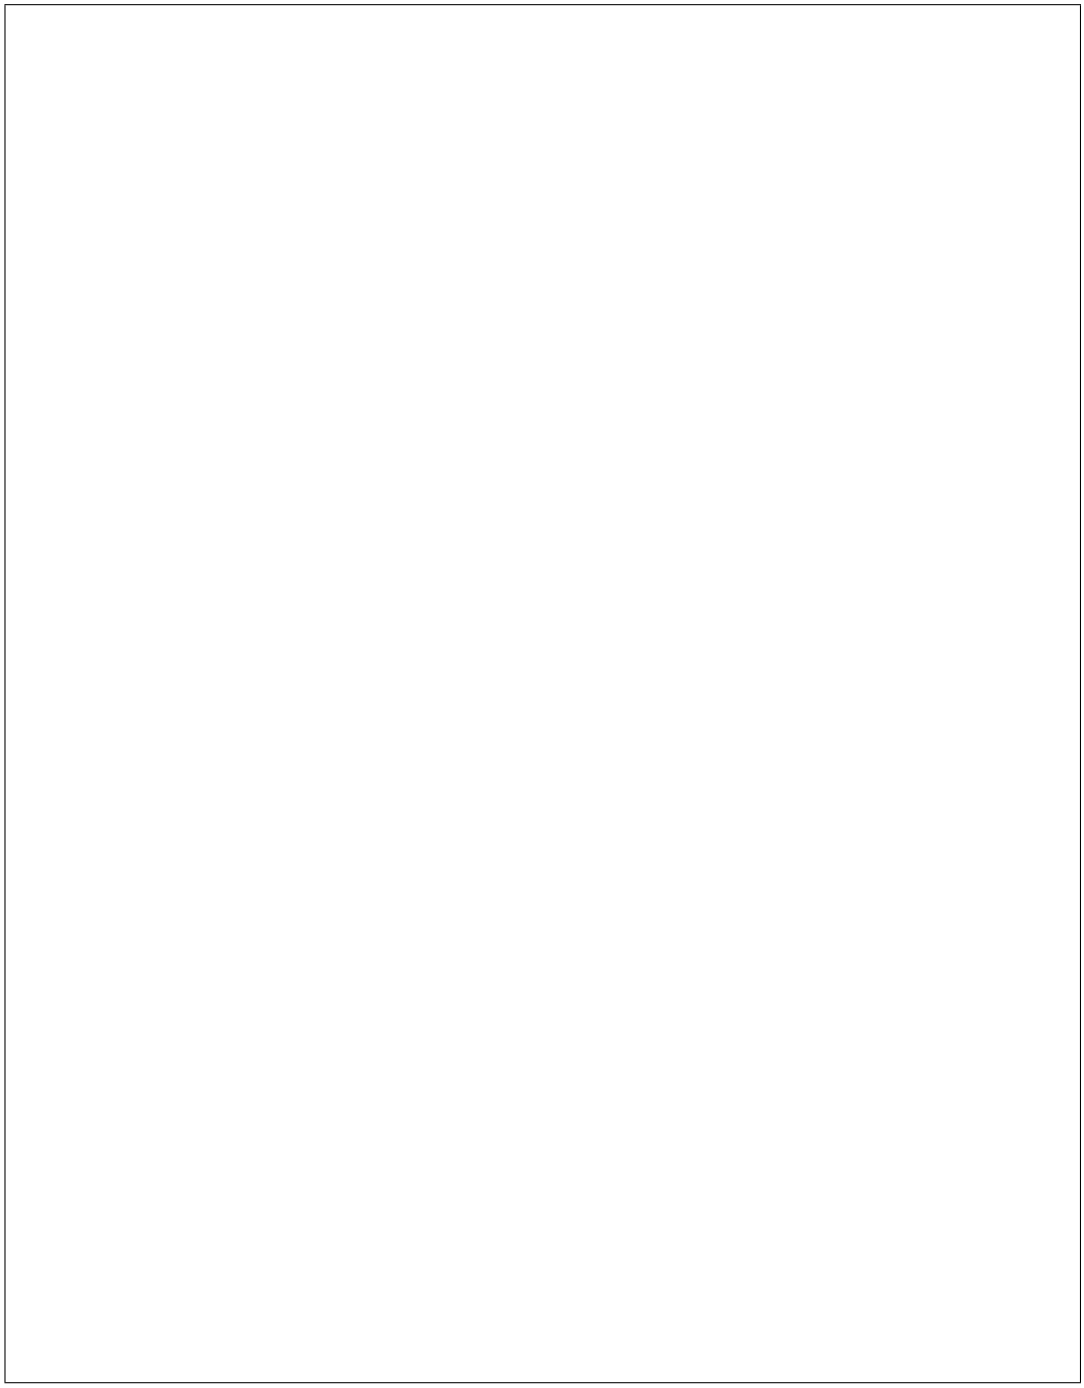

Open region: every point in the region is the centre of some disc lying entirely in the region (ie. an open region doesn't include the boundary points).

Connected region: Any two points in D can be joined by a path lying entirely in D.

Another interesting result is that if $\int_C \boldsymbol{F} \cdot \boldsymbol{dr}$ is path independent in some region D, then $\oint_{C'} \boldsymbol{F} \cdot \boldsymbol{dr} = 0$ for *every* closed path C' in D. Here the symbol "$\oint$" indicates the integral is over a *closed* curve.

Proof:

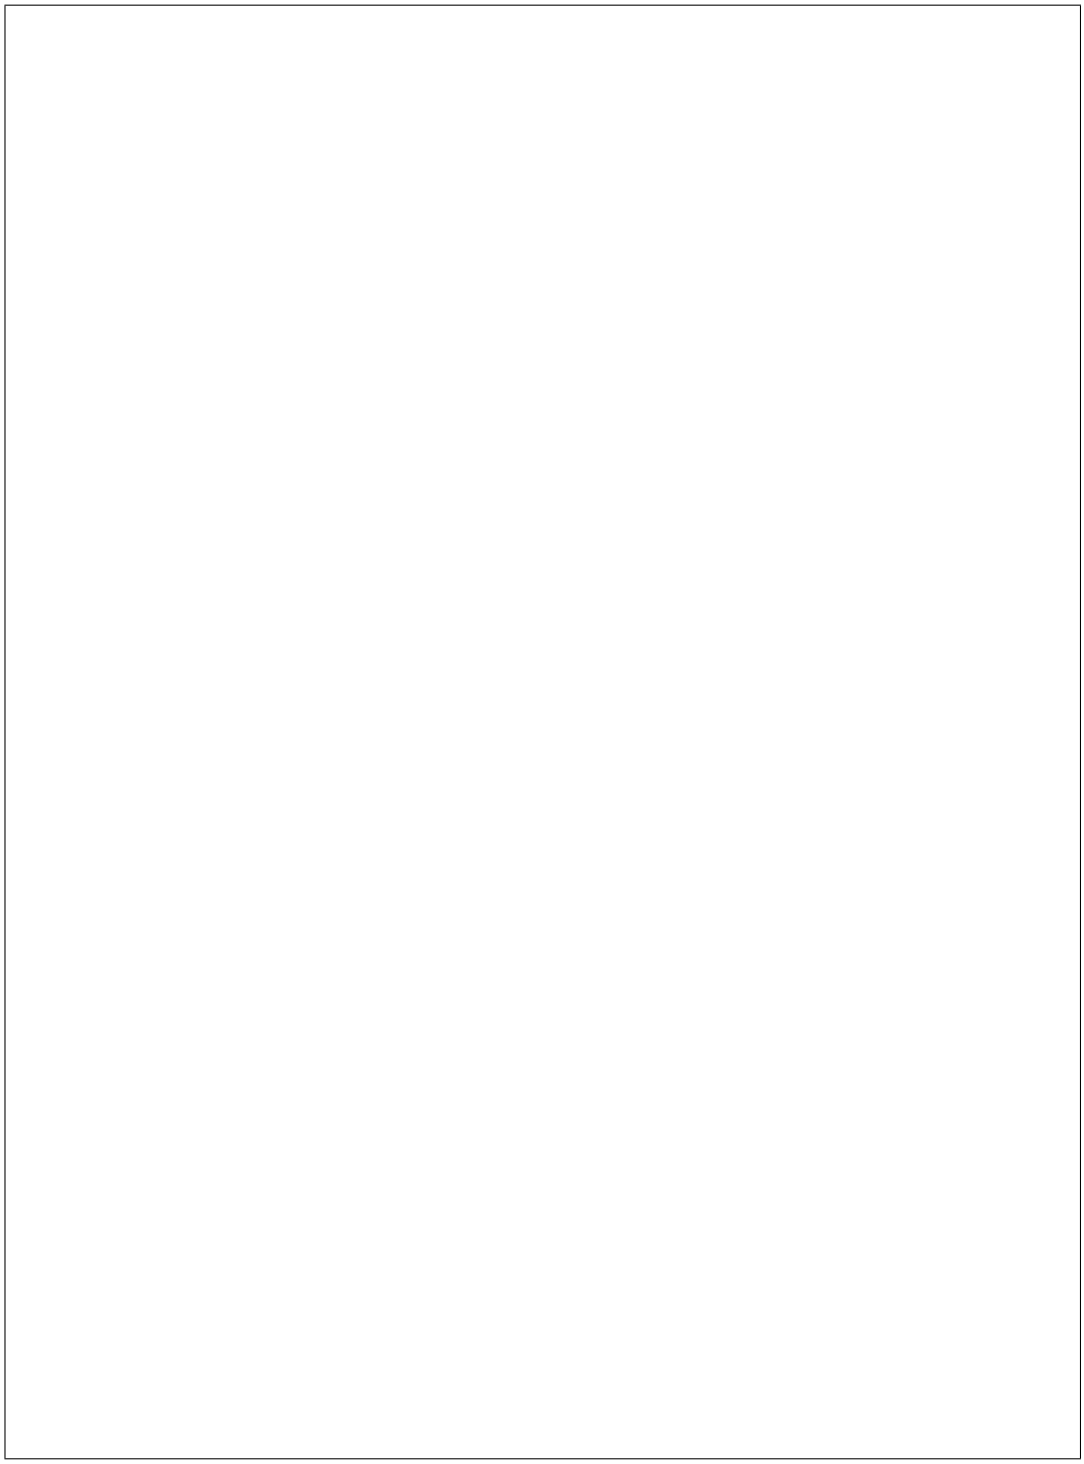

Perhaps it is not surprising that the converse is also true. That is, if $\oint_{C'} \boldsymbol{F} \cdot \boldsymbol{dr} = 0$ for *every* closed path C' in some region D, then $\int_C \boldsymbol{F} \cdot \boldsymbol{dr}$ is path independent in D.

Proof:

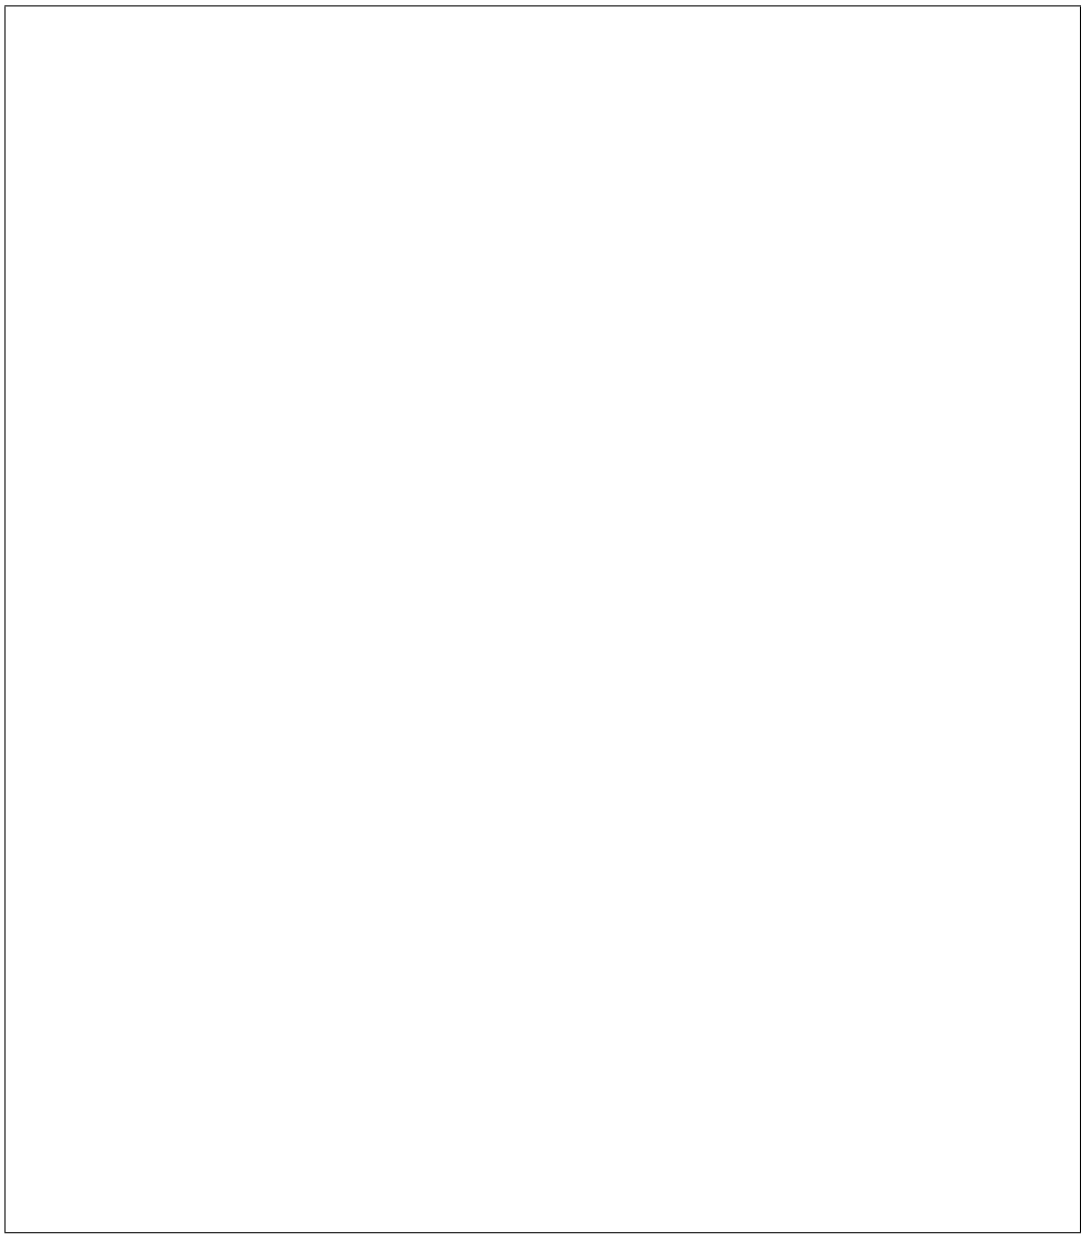

We are looking at these results carefully because we ultimately want a simple way of checking whether or not a vector field is conservative. We are not quite there yet, but in the next section, we will arrive at a surprisingly simple test for a conservative vector field.

Note also that more details of these proofs (with slightly more mathematical rigour) can be found in Stewart, pages 1099 – 1103.

20 Green's theorem and a test for conservative fields

By the end of this section, you should be able to answer the following questions:

- What is Green's theorem and under what conditions can it be applied?
- How do you apply Green's theorem?
- Given a vector field in two dimensions, how can we test whether or not it is conservative?

20.1 The story so far

The following diagram summarises the relationships between conservative vector fields, path independent line integrals and closed line integrals we have seen so far.

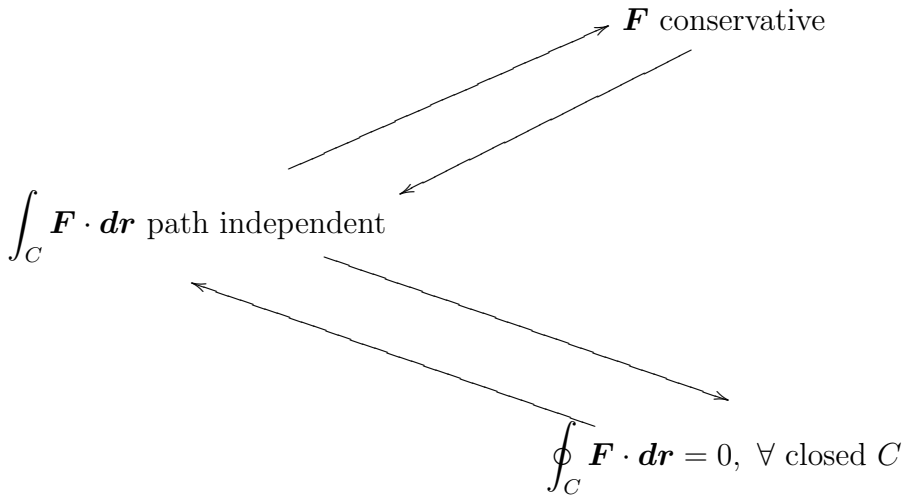

20.2 Clairaut's theorem and consequences

Suppose a function of two variables f is defined on a disc D that contains the point (a, b). If the functions $\dfrac{\partial^2 f}{\partial x \partial y}$ and $\dfrac{\partial^2 f}{\partial y \partial x}$ are both continuous on D, then

$$\frac{\partial^2 f}{\partial x \partial y}(a, b) = \frac{\partial^2 f}{\partial y \partial x}(a, b).$$

Say we have a conservative vector field $\boldsymbol{F} = F_1\boldsymbol{i} + F_2\boldsymbol{j}$. This means that there exists an $f(x, y)$ such that

$$F_1 = \frac{\partial f}{\partial x}, \quad F_2 = \frac{\partial f}{\partial y}.$$

An immediate consequence of Clairaut's theorem is that

$$\frac{\partial F_1}{\partial y} = \frac{\partial^2 f}{\partial y \partial x} = \frac{\partial^2 f}{\partial x \partial y} = \frac{\partial F_2}{\partial x}.$$

In otherwords, we have the following:

If $\boldsymbol{F} = F_1\boldsymbol{i} + F_2\boldsymbol{j}$ is a conservative vector field, then

$$\frac{\partial F_1}{\partial y} = \frac{\partial F_2}{\partial x}.$$

Let's add this to our diagram:

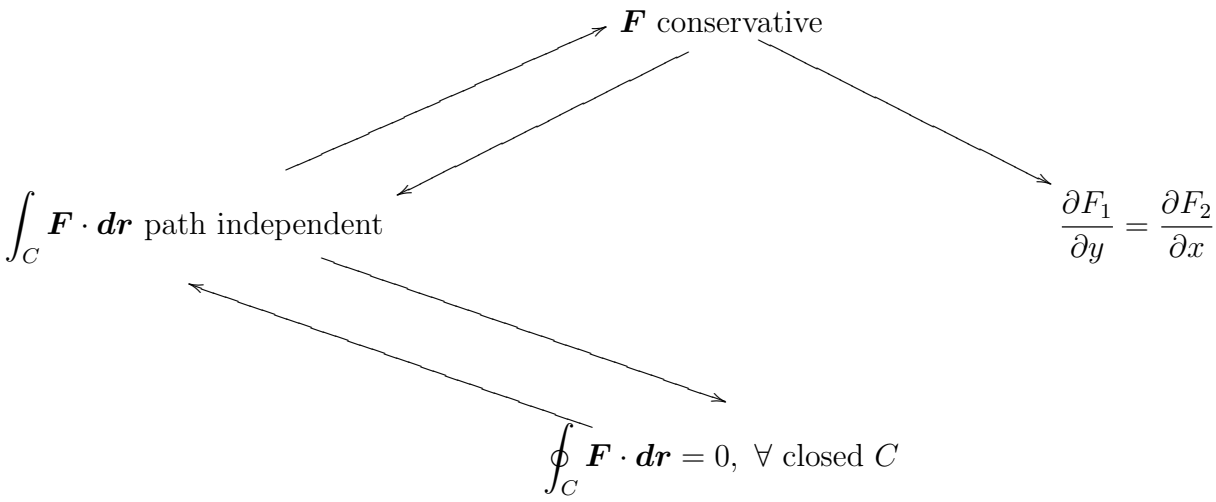

If we can reverse the new arrow, then we would have the criterion that we need! That is, the condition

$$\frac{\partial F_1}{\partial y} = \frac{\partial F_2}{\partial x}$$

would be a test for a conservative vector field. To do this, we require one more piece of the puzzle. That is Green's theorem.

20.3 Green's theorem

Let D be a region in the x-y plane bounded by a piecewise-smooth, simple closed curve C, which is traversed with D always on the left. Let $F_1(x,y)$, $F_2(x,y)$, $\dfrac{\partial F_1}{\partial y}$ and $\dfrac{\partial F_2}{\partial x}$ be continuous in D. Then

$$\iint_D \left(\frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y}\right)\, dx\, dy = \oint_C (F_1 dx + F_2 dy).$$

This theorem relates a double integral to a line integral over a closed curve. For example, we can use Green's theorem to evaluate complicated line integrals by treating them as double integrals, or vice versa.

Regarding our discussion on conservative vector fields, we have the following *corollary* to Green's theorem:

$$\text{If } \frac{\partial F_1}{\partial y} = \frac{\partial F_2}{\partial x}, \text{ then } \oint_C \boldsymbol{F}\cdot\boldsymbol{dr} = 0.$$

Note that $\boldsymbol{F} = F_1\boldsymbol{i} + F_2\boldsymbol{j}$.

If we add this to our diagram, we can now link any four statements via the arrows. In otherwords all four statements are equivalent.

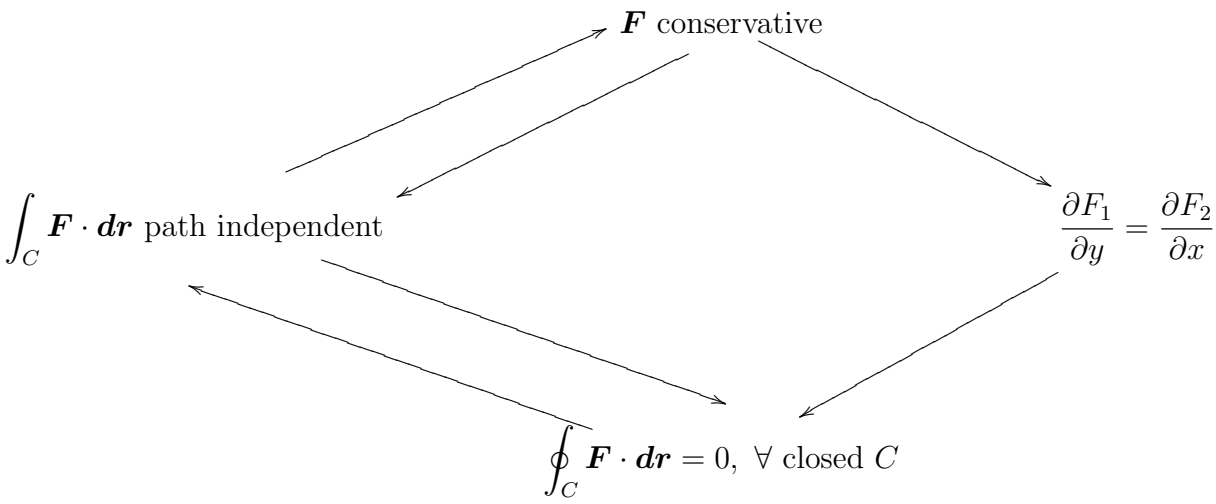

In particular, we now have a test to determine whether or not a given two dimensional vector field is conservative:

$$\text{The vector field } \boldsymbol{F} \text{ is conservative if and only if } \frac{\partial F_1}{\partial y} = \frac{\partial F_2}{\partial x}.$$

20.3.1 Find the work done by the force $\boldsymbol{F} = x^2y\boldsymbol{i} + xy^2\boldsymbol{j}$ anticlockwise around the circle with centre at the origin and radius a.

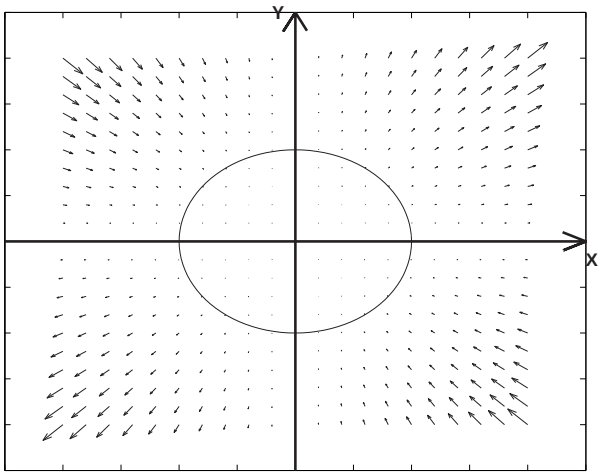

20.3.2 **Evaluate the line integral** $\int_C 2xy\ dx + (x^2 + 3y^2)\ dy$**, where** C **is the path from** $(0, 1)$ **to** $(1, 0)$ **along** $y = (x - 1)^2$ **and then from** $(1, 0)$ **to** $(2, 1)$ **along** $y = x - 1$**.**

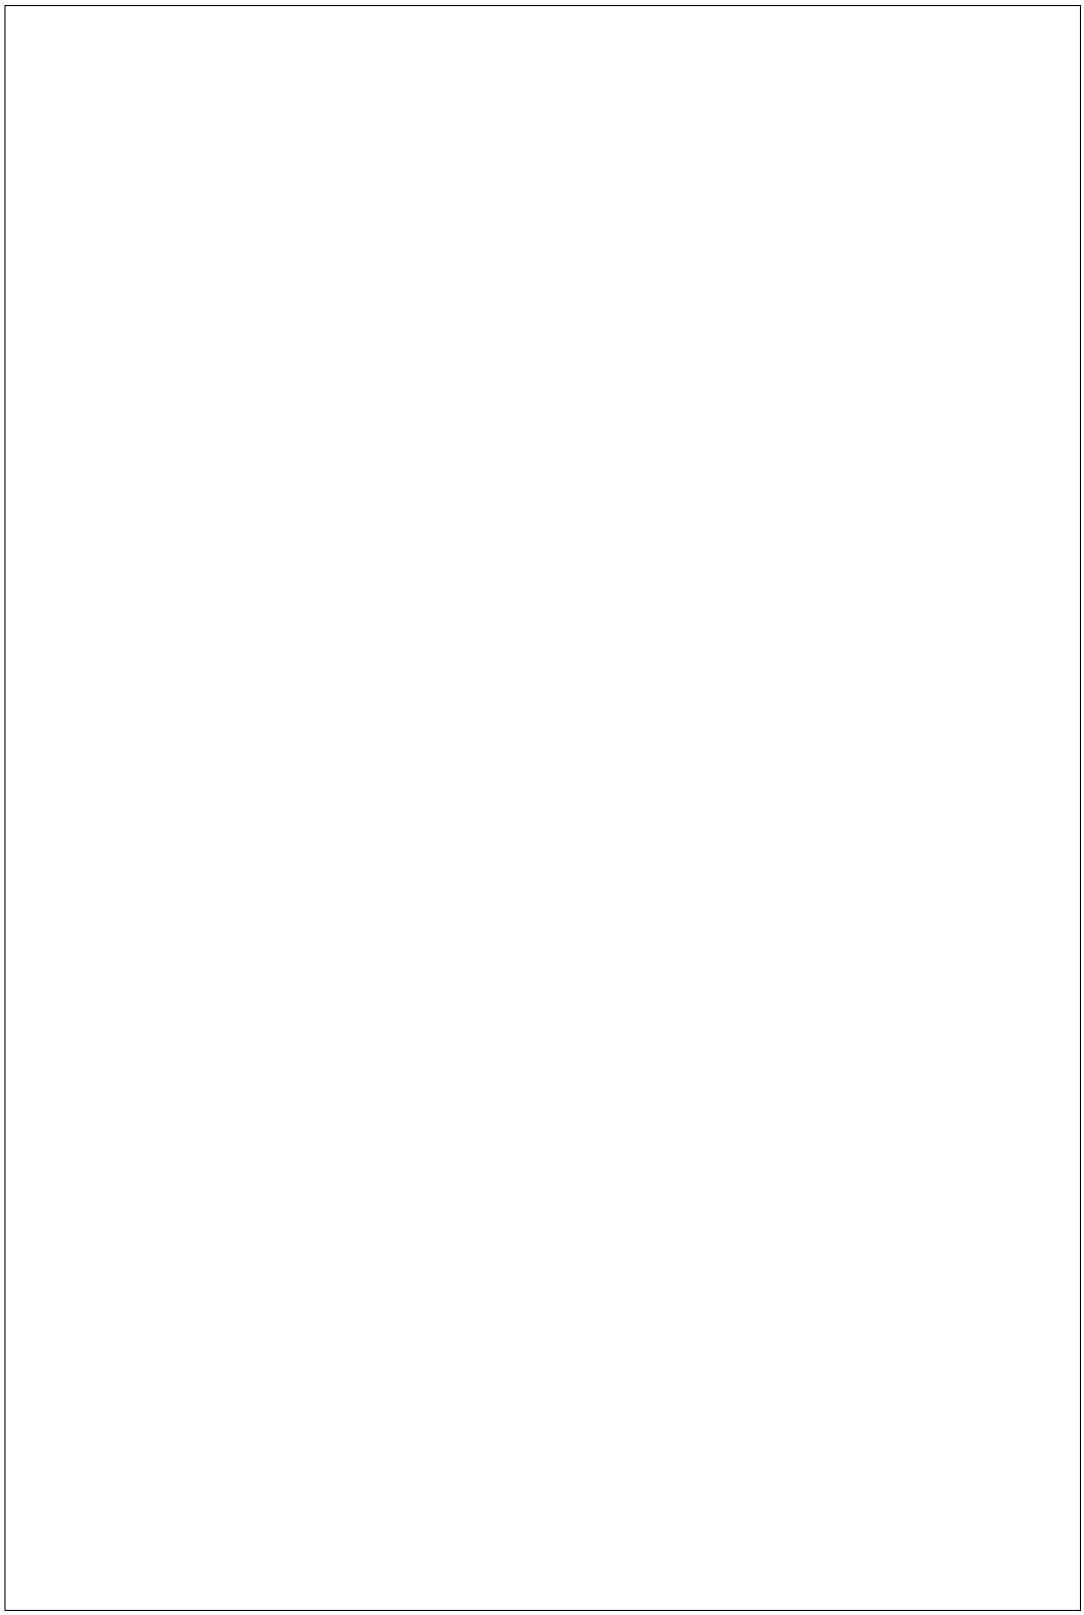

20.3.3 Evaluate $\int_C (3+2xy)dx + (x^2-3y^2)dy$ where C is the curve parametrised by $\boldsymbol{r}(t) = (1-\cos(\pi t))\boldsymbol{i} + (1+\sin^3(\pi t))\boldsymbol{j}$ for $0 \leq t \leq 1/2$.

21 Flux of a vector field

By the end of this section, you should be able to answer the following questions:

- What is the flux of a constant vector field across a flat surface in 3D?
- What is the flux of a vector field across a plane curve in 2D?

In this section we introduce the concept of *flux*: In three dimensions, the flux of a vector field across a given surface is defined to be the "flow rate" of the vector field through the surface.

Since many vector fields involve no motion (eg. electric fields, magnetic fields), this definition can be very difficult to comprehend at first. A nice context for working with flux in order to understand its definition is by considering the velocity vector of a fluid (so now we do have motion). In three dimensions, the flux of a fluid across a surface is given in units of volume per unit time. In other words, the flux tells us how much of the fluid (volume) passes through a given surface in one second.

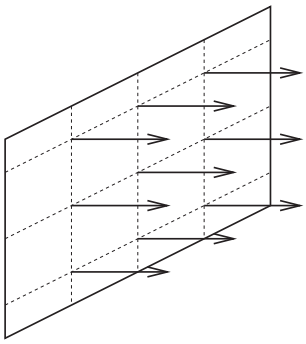

Consider a river flowing at a constant velocity of $2m/s$ in only one direction. Now imagine placing a $3m$ square fishing net into the river so that it somehow stays perpendicular to the flow of the river. What is the flux of the water through the net?

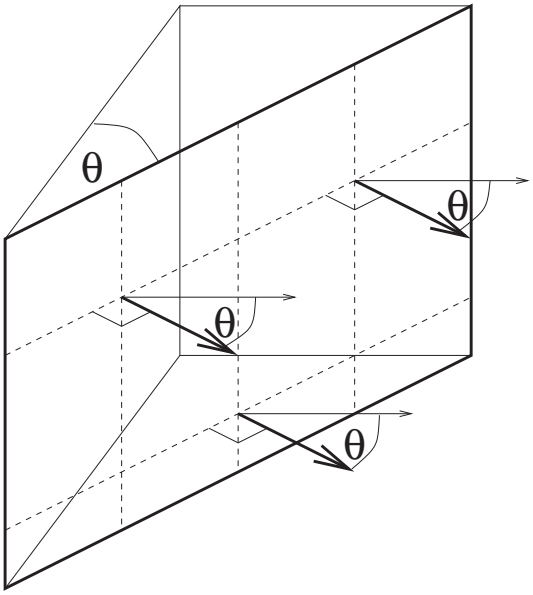

Now if we rotate the net through an angle θ, what is the flux through the net?

21.1 Flux in 2D

Before we look at the flux of a vector field through more general surfaces, let's look at flux in two dimensions, by considering the flow of a two dimensional fluid through a curve in the x-y plane. Note that in this context of a fluid in 2D, flux has dimensions *area* per unit time.

To start, consider the problem of calculating the flux of a fluid with constant velocity $\boldsymbol{v} = 2\boldsymbol{i}$ through a line segment C perpendicular to the flow, where C is given by

$$C = \{(x, y) \mid x = 2,\ 2 \leq y \leq 6\}.$$

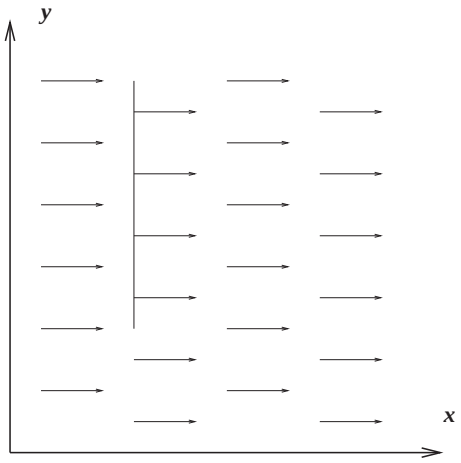

Now consider calculating the flux of the velocity vector $\boldsymbol{v}(x, y)$ in the x-y plane through a curve C.

We first divide C up into arcs of length ΔS, and *approximate* $\boldsymbol{v}$ as constant over each arc.

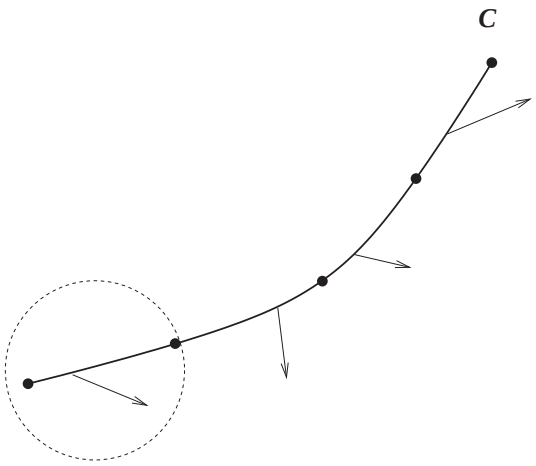

This constant vector over each arc shall be evaluated at a representative point in each arc, say $P^* = (x^*, y^*)$. We also approximate the arc as a straight line, so that

$$\Delta S \approx \sqrt{(\Delta x)^2 + (\Delta y)^2} \approx |\boldsymbol{r}'(t)|\ \Delta t.$$

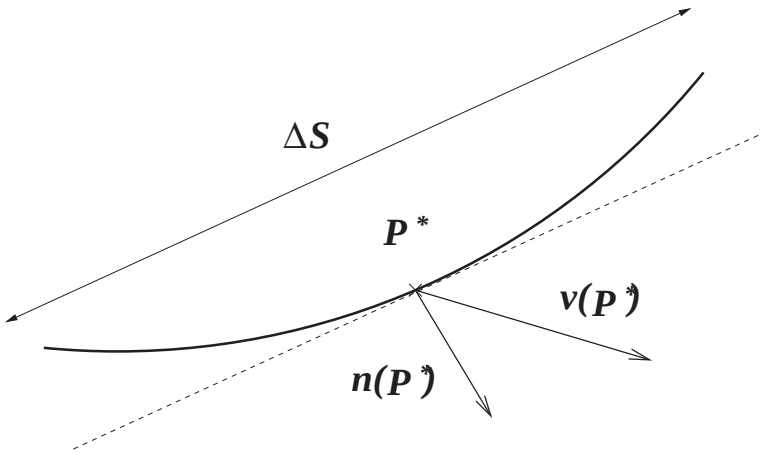

The component of $\boldsymbol{v}$ which is perpendicular to C (over ΔS) is $\approx \boldsymbol{v}(P^*) \cdot \boldsymbol{n}(P^*)$. We then have

$$\begin{aligned} &\text{flux through one arc } \approx \boldsymbol{v}(P^*) \cdot \boldsymbol{n}(P^*)\Delta S. \\ \Rightarrow\ &\text{total flux through } C \approx \sum_i \boldsymbol{v}(P_i^*) \cdot \boldsymbol{n}(P_i^*)\Delta S_i. \end{aligned}$$

If we take the limit as $\Delta S \to 0$, we obtain an exact expression for the flux over the entire curve C as a line integral:

$$\text{Flux } = \int_C \boldsymbol{v} \cdot \boldsymbol{n} \; dS,$$

where $\boldsymbol{n}$ is a unit vector normal to C.

We use this expression as a *definition* of flux of any two dimensional vector field $\boldsymbol{v}$ across a plane curve C. Note then that

$$\text{dimensions of flux (in 2D) } = (\text{dimensions of } \boldsymbol{v}) \; \times \; (\text{distance}).$$

21.1.1 Evaluating flux in 2D

To evaluate the line integral in the definition of flux, we need a parametrisation of C, say $\boldsymbol{r}(t) = x(t)\boldsymbol{i} + y(t)\boldsymbol{j}$ for $a \leq t \leq b$ (say). We define

$$\boldsymbol{r}'(t) = \dot{x} \; \boldsymbol{i} + \dot{y} \; \boldsymbol{j}.$$

A unit tangent vector to C is then given by

$$\boldsymbol{T} = \frac{\boldsymbol{r}'(t)}{|\boldsymbol{r}'(t)|}.$$

By the definition of vector cross product, and since $\boldsymbol{k}$ is a unit vector normal to the x-y plane, being careful of the direction of $\boldsymbol{n}$, we can take

$$\boldsymbol{n} = \boldsymbol{T} \times \boldsymbol{k} = \frac{1}{|\boldsymbol{r}'(t)|} \begin{vmatrix} \boldsymbol{i} & \boldsymbol{j} & \boldsymbol{k} \\ \dot{x} & \dot{y} & 0 \\ 0 & 0 & 1 \end{vmatrix}$$

$$\begin{aligned} &= \frac{1}{|\boldsymbol{r}'(t)|}(\dot{y} \; \boldsymbol{i} - \dot{x} \; \boldsymbol{j}) \\ \Rightarrow \; \boldsymbol{v} \cdot \boldsymbol{n} &= \frac{\boldsymbol{v} \cdot (\dot{y} \; \boldsymbol{i} - \dot{x} \; \boldsymbol{j})}{|\boldsymbol{r}'(t)|} \\ &= \frac{v_1\dot{y} - v_2\dot{x}}{|\boldsymbol{r}'(t)|}, \end{aligned}$$

where $\boldsymbol{v}(x, y) = v_1(x, y)\boldsymbol{i} + v_2(x, y)\boldsymbol{j}$. Noting also that in the integral we have $dS = |\boldsymbol{r}'(t)| \; dt$, we then have a means of evaluating the line integral (2D flux integral) as

$$\int_C \boldsymbol{v} \cdot \boldsymbol{n} \; dS = \int_{t=a}^{t=b} (v_1(t)\dot{y} - v_2(t)\dot{x}) \; dt.$$

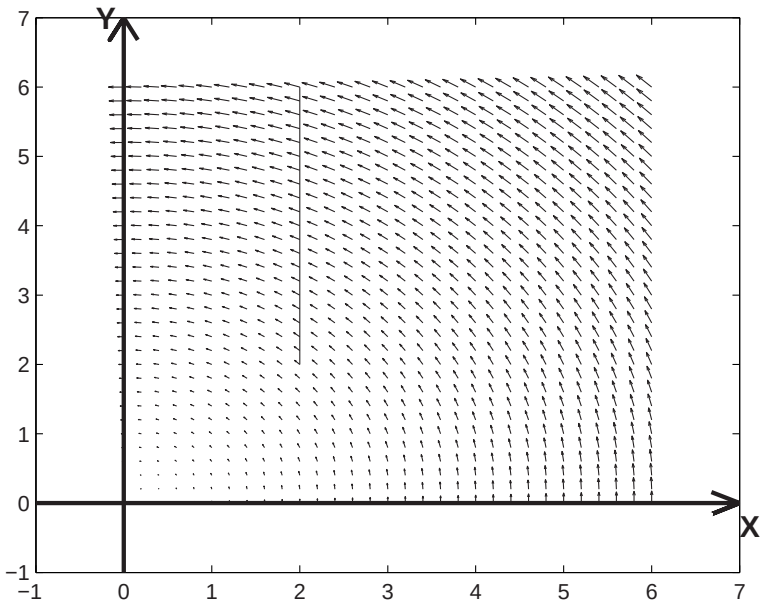

21.1.2 Calculate the flux of $\boldsymbol{v} = -y\boldsymbol{i} + x\boldsymbol{j}$ (in the positive x direction) across the line $x = 2$ (for $2 \leq y \leq 6$).

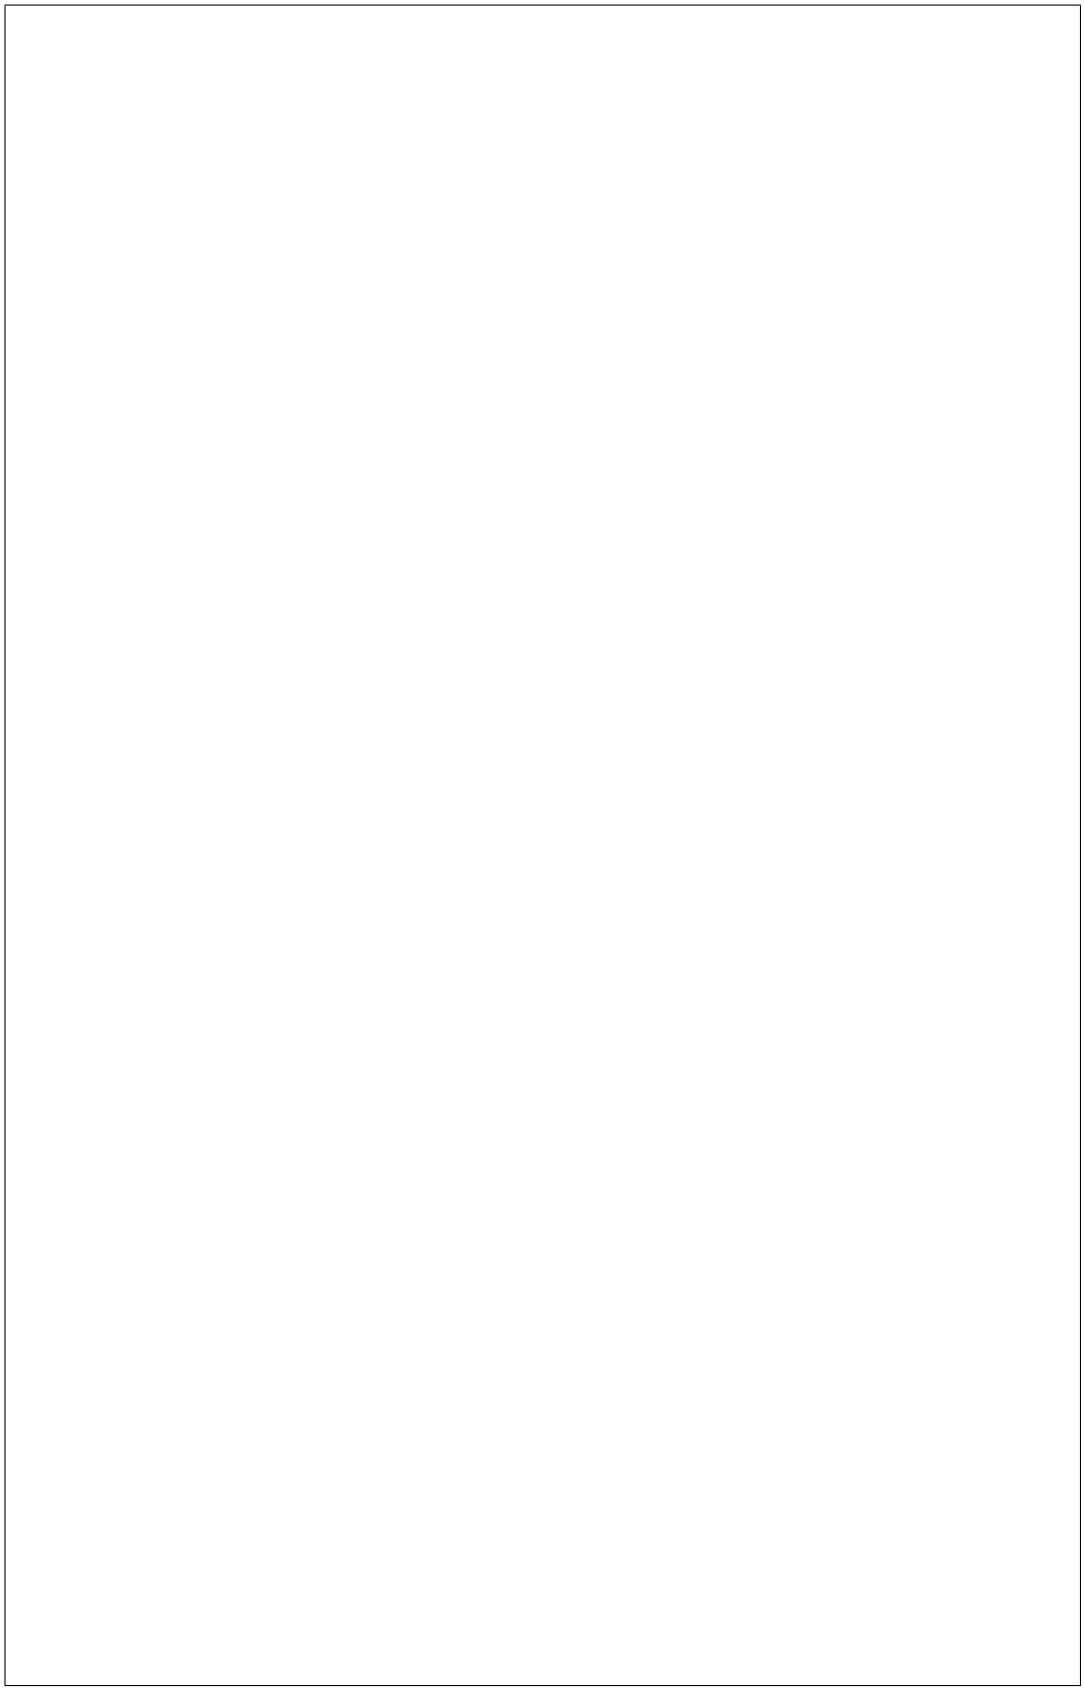

21.2 Outward flux across a closed curve in the plane

Let C be a piecewise-smooth, simple closed curve. Let $v_1(x, y)$, $v_2(x, y)$ be continuous in the region bounded by C. (Note that these are some of the conditions of Green's theorem!)

The net outward flux of $\boldsymbol{v} = v_1\boldsymbol{i} + v_2\boldsymbol{j}$ across C is given by

$$\text{Net outward flux } = \oint_C \boldsymbol{v} \cdot \boldsymbol{n}\, dS,$$

where $\boldsymbol{n}$ is a unit vector normal to C, directed outward from the region bounded by C.

21.2.1 Calculate the outward flux of $\boldsymbol{v} = xy\boldsymbol{i} + xy\boldsymbol{j}$ across the curve from (2,0) to (-2,0) via the semicircle of radius 2 centred at the origin (for $y \geq 0$) followed by the straight line from (-2,0) to (2,0).

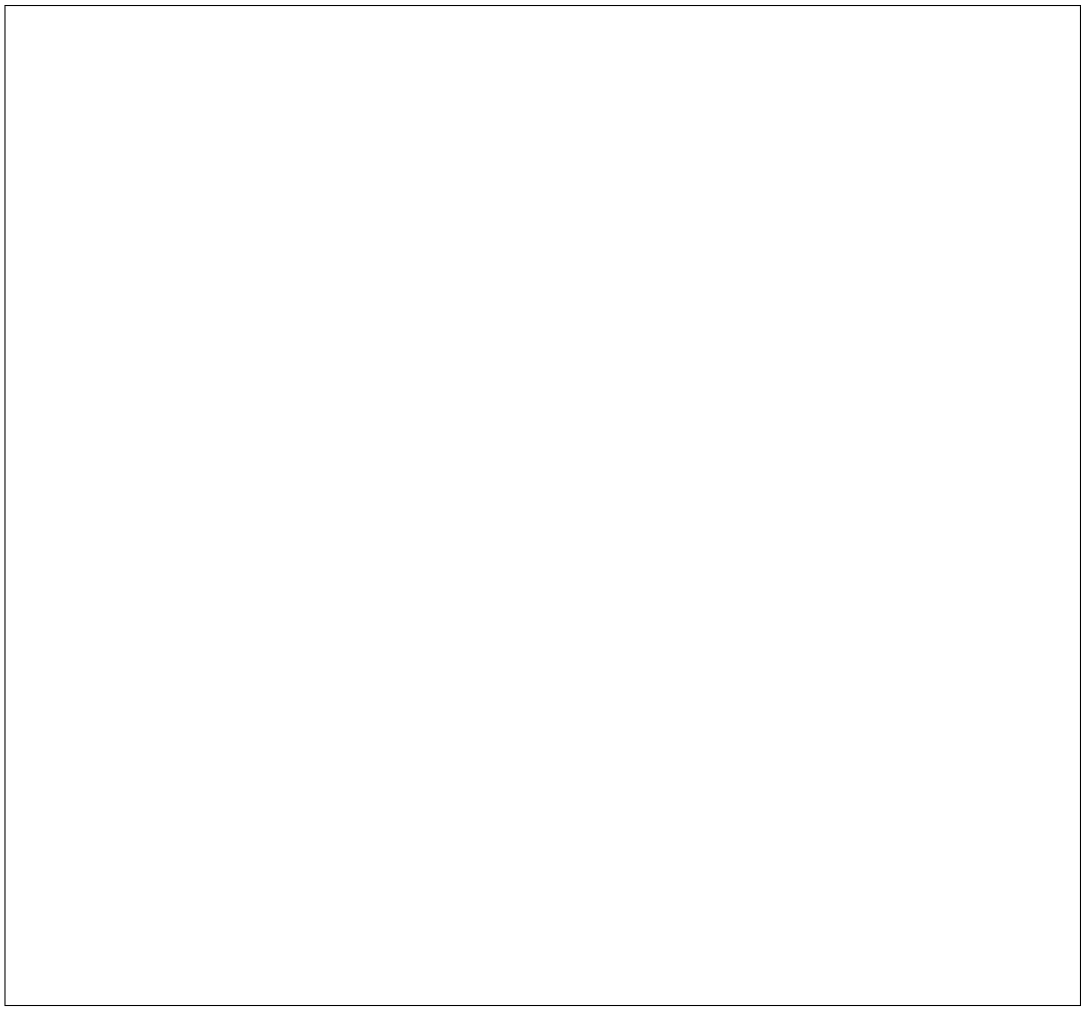

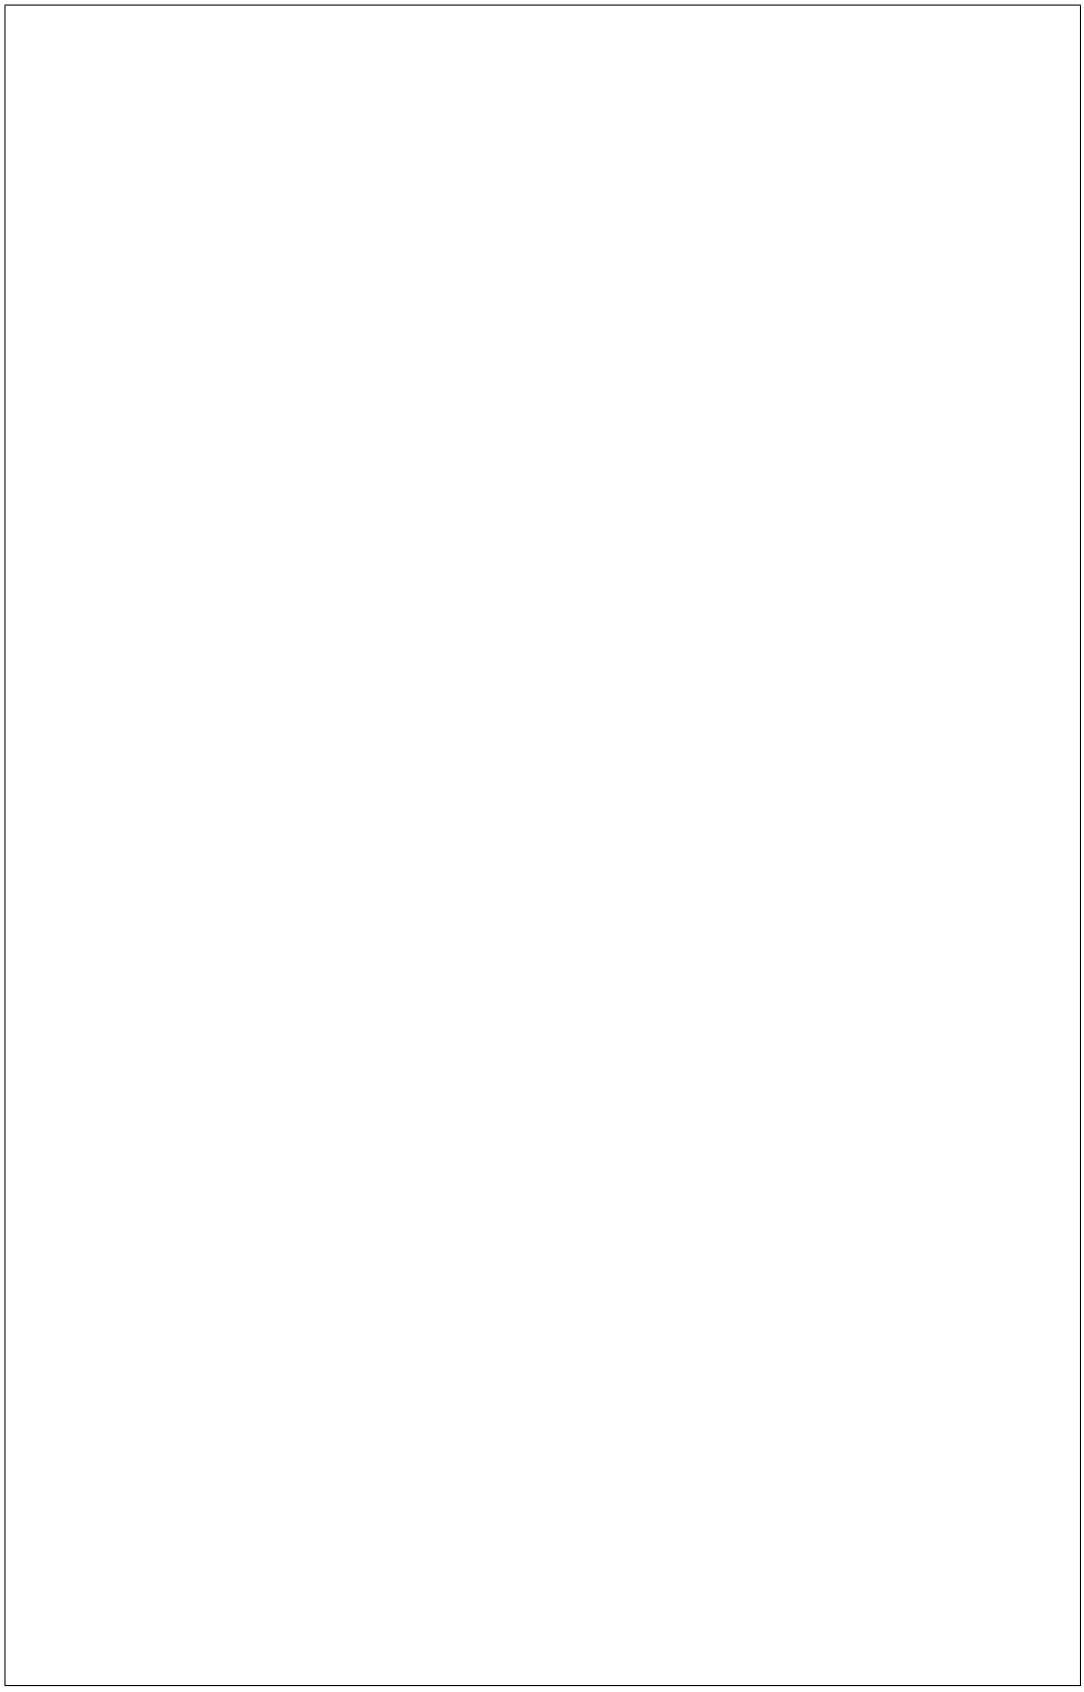

22 Divergence of a vector field (div)

By the end of this section, you should be able to answer the following questions:

- How do you calculate the divergence of a given vector field?
- What is the significance of divergence?
- How does it relate to flux?

In this section we introduce the concept of *divergence* of a vector field.

22.1 Calculating divergence

Let

$$\boldsymbol{v}(x,y,z) = v_1(x,y,z)\boldsymbol{i} + v_2(x,y,z)\boldsymbol{j} + v_3(x,y,z)\boldsymbol{k}$$

be a differentiable vector function. Then the function

$$\text{div } \boldsymbol{v} = \frac{\partial v_1}{\partial x} + \frac{\partial v_2}{\partial y} + \frac{\partial v_3}{\partial z} = \nabla \cdot \boldsymbol{v}$$

is called the divergence of $\boldsymbol{v}$. Note div $\boldsymbol{v}$ is a scalar quantity.

Divergence has an analogous definition in two dimensions. For

$$\boldsymbol{F}(x,y) = F_1(x,y)\boldsymbol{i} + F_2(x,y)\boldsymbol{j} \quad \Rightarrow \quad \text{div}\boldsymbol{F} = \frac{\partial F_1}{\partial x} + \frac{\partial F_2}{\partial y}.$$

22.1.1 Example: $\boldsymbol{v} = xy^2\boldsymbol{i} + xyz\boldsymbol{j} + yz^2\boldsymbol{k}$. Find div $\boldsymbol{v}$

22.2 Understanding div in two dimensions.

Consider the flow of a two dimensional fluid with continuous velocity field $\boldsymbol{v}(x, y) = v_1(x, y)\boldsymbol{i} + v_2(x, y)\boldsymbol{j}$. Our aim is to calculate the outward flux from a small rectangle in the plane of area $\Delta x \Delta y$ as in the diagram below.

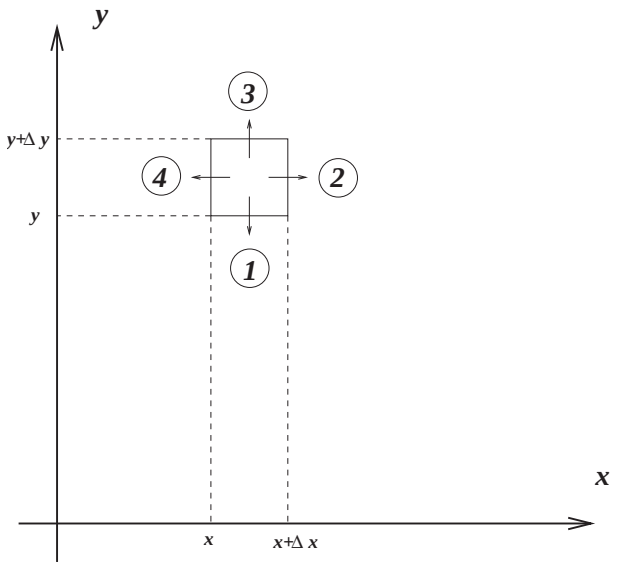

We first approximate the flux across each of the four sides of the rectangle. In each case the approximation will be $\boldsymbol{v} \cdot \boldsymbol{n} \Delta S$, where we assume $\boldsymbol{v}$ is constant over each edge. Also let $x^* \in [x, x + \Delta x]$ and $y^* \in [y, y + \Delta y]$ represent chosen points in each interval.

Edge 1: we evaluate $\boldsymbol{v}$ at (x^*, y) and assume it is constant across the entire edge. An outwardly pointing unit normal vector is $-\boldsymbol{j}$.

$$\text{flux} \approx \boldsymbol{v}(x^*, y) \cdot (-\boldsymbol{j})\ \Delta x.$$

Edge 2: we evaluate $\boldsymbol{v}$ at $(x + \Delta x, y^*)$ and assume it is constant across the entire edge. An outwardly pointing unit normal vector is $\boldsymbol{i}$.

$$\text{flux} \approx \boldsymbol{v}(x + \Delta x, y^*) \cdot (\boldsymbol{i})\ \Delta y.$$

Edge 3: we evaluate $\boldsymbol{v}$ at $(x^*, y + \Delta y)$ and assume it is constant across the entire edge. An outwardly pointing unit normal vector is $\boldsymbol{j}$.

$$\text{flux} \approx \boldsymbol{v}(x^*, y + \Delta y) \cdot (\boldsymbol{j})\ \Delta x.$$

Edge 4: we evaluate $\boldsymbol{v}$ at (x, y^*) and assume it is constant across the entire edge. An outwardly pointing unit normal vector is $-\boldsymbol{i}$.

$$\text{flux} \approx \boldsymbol{v}(x, y^*) \cdot (-\boldsymbol{i})\ \Delta y.$$

Combining all four terms gives an approximation to the net outward flux:

net outward flux

$$\begin{aligned}
&\approx (\boldsymbol{v}(x+\Delta x, y^*) - \boldsymbol{v}(x, y^*)) \cdot \boldsymbol{i}\ \Delta y + (\boldsymbol{v}(x^*, y+\Delta y) - \boldsymbol{v}(x^*, y)) \cdot \boldsymbol{j}\ \Delta x \\
&= \left(\frac{\boldsymbol{v}(x+\Delta x, y^*) - \boldsymbol{v}(x, y^*)}{\Delta x}\right) \cdot \boldsymbol{i}\Delta x\Delta y + \left(\frac{\boldsymbol{v}(x^*, y+\Delta y) - \boldsymbol{v}(x^*, y)}{\Delta y}\right) \cdot \boldsymbol{j}\Delta x\Delta y \\
&= \left(\frac{v_1(x+\Delta x, y^*) - v_1(x, y^*)}{\Delta x} + \frac{v_2(x^*, y+\Delta y) - v_2(x^*, y)}{\Delta y}\right) \Delta x\Delta y \\
&\approx \left(\frac{\partial v_1}{\partial x} + \frac{\partial v_2}{\partial y}\right) \Delta x\Delta y. \\
&= \operatorname{div}(\boldsymbol{v})\Delta x\Delta y.
\end{aligned}$$

Hence, we have

$$\frac{\text{flux out of a rectangle}}{\text{area of rectangle}} \approx \operatorname{div}(\boldsymbol{v}).$$

If we take the limit as the dimensions of the rectangle appproach 0, we have

$$\operatorname{div}(\boldsymbol{v}) = \lim_{\Delta A \to 0} \frac{\text{flux out of } \Delta A}{\Delta A}.$$

In other words, $\operatorname{div}(\boldsymbol{v})$ is the "outward flux density" of $\boldsymbol{v}$ at a given point.

This concept generalises quite naturally to three dimensions:

$$\operatorname{div}(\boldsymbol{v}(x, y, z)) = \lim_{\Delta V \to 0} \frac{\text{flux out of } \Delta V}{\Delta V}.$$

In the context of fluids (our main focus so far) we can say $\operatorname{div}(\boldsymbol{v}(x, y, z))$ measures the tendency of the fluid to "diverge" from the point (x, y, z).

22.3 Outward flux across a closed curve in the plane (revisited)

One final calculation uses the divergence to calculate the net outward flux of $\boldsymbol{v}$ across a closed curve. We have already seen that we can evaluate this quantity by calculating $\oint_C \boldsymbol{v} \cdot \boldsymbol{n}\, dS$.

Now let D be a region in the x-y plane bounded by a piecewise-smooth, simple closed curve C, which is traversed with D always on the left. Let $v_1(x, y)$, $v_2(x, y)$ have continuous derivatives in D (again the conditions of Green's theorem!).

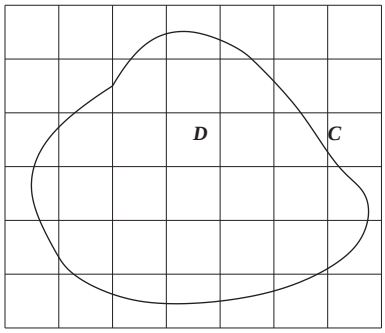

By the previous calculation involving divergence, we can also approximate the outward flux from the region by dividing D up into small rectangles and approximating the net outward flux across each rectangle. We know that for one rectangle,

$$\text{outward flux of one rectangle } \approx \operatorname{div}(\boldsymbol{v}(x^*, y^*))\Delta x \Delta y,$$

where (x^*, y^*) is some point inside the rectangle. We repeat this for each rectangle containing part of the region D, so that

$$\text{net outward flux across } C \approx \sum \operatorname{div}(\boldsymbol{v}(x^*, y^*))\Delta x \Delta y.$$

Taking the limit as $\Delta x, \Delta y \to 0$, we have

$$\text{net outward flux across } C = \iint_D \operatorname{div}(\boldsymbol{v}(x, y))\, dA,$$

the double integral of the region D.

To obtain the flux, we integrate the flux density over the region. Compare this with the context of mass density: to obtain the mass, we integrate the mass density over the region.

Finally, the two ways of calculating the same quantity must obviously be equal:

$$\oint_C \boldsymbol{v}(x, y) \cdot \boldsymbol{n}\, dS = \iint_D \operatorname{div}(\boldsymbol{v}(x, y))\, dA.$$

22.4 Relationship to Green's theorem

We have seen how to evaluate the 2D flux integral:

$$\oint_C \boldsymbol{v} \cdot \boldsymbol{n} \; dS = \int_{t=a}^{t=b} (v_1(t)\dot{y} - v_2(t)\dot{x}) \; dt.$$

This can be rewritten as

$$\oint_C \boldsymbol{v} \cdot \boldsymbol{n} \; dS = \oint_C v_1 \; dy - v_2 \; dx.$$

If we define $F_1(x, y) = -v_2(x, y)$ and $F_2(x, y) = v_1(x, y)$, we then have

$$\oint_C \boldsymbol{v} \cdot \boldsymbol{n} \; dS = \oint_C F_1 \; dx + F_2 \; dy.$$

We also have

$$\text{div}(\boldsymbol{v}) = \frac{\partial v_1}{\partial x} + \frac{\partial v_2}{\partial y} = \frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y},$$

so that

$$\iint_D \text{div}(\boldsymbol{v}) \; dA = \iint_D \left(\frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y} \right) \; dA.$$

This tells us that in terms of the new vector field

$$\boldsymbol{F} = -v_2\boldsymbol{i} + v_1\boldsymbol{j} = F_1\boldsymbol{i} + F_2\boldsymbol{j},$$

the two ways of calculating flux are given by

$$\iint_D \left(\frac{\partial F_2}{\partial x} - \frac{\partial F_1}{\partial y} \right) \; dx \; dy = \oint_C (F_1 dx + F_2 dy).$$

This is none other than Green's theorem. So the flux identity we obtained at the bottom of the previous page is just Green's theorem in disguise. We shall call this the *flux form* of Green's theorem.

22.4.1 Use the flux form of Green's theorem to calculate the outward flux of $\boldsymbol{v} = xy\boldsymbol{i} + xy\boldsymbol{j}$ across the curve from (2,0) to (-2,0) via the semicircle of radius 2 centred at the origin (for $y \geq 0$) followed by the straight line from (-2,0) to (2,0).

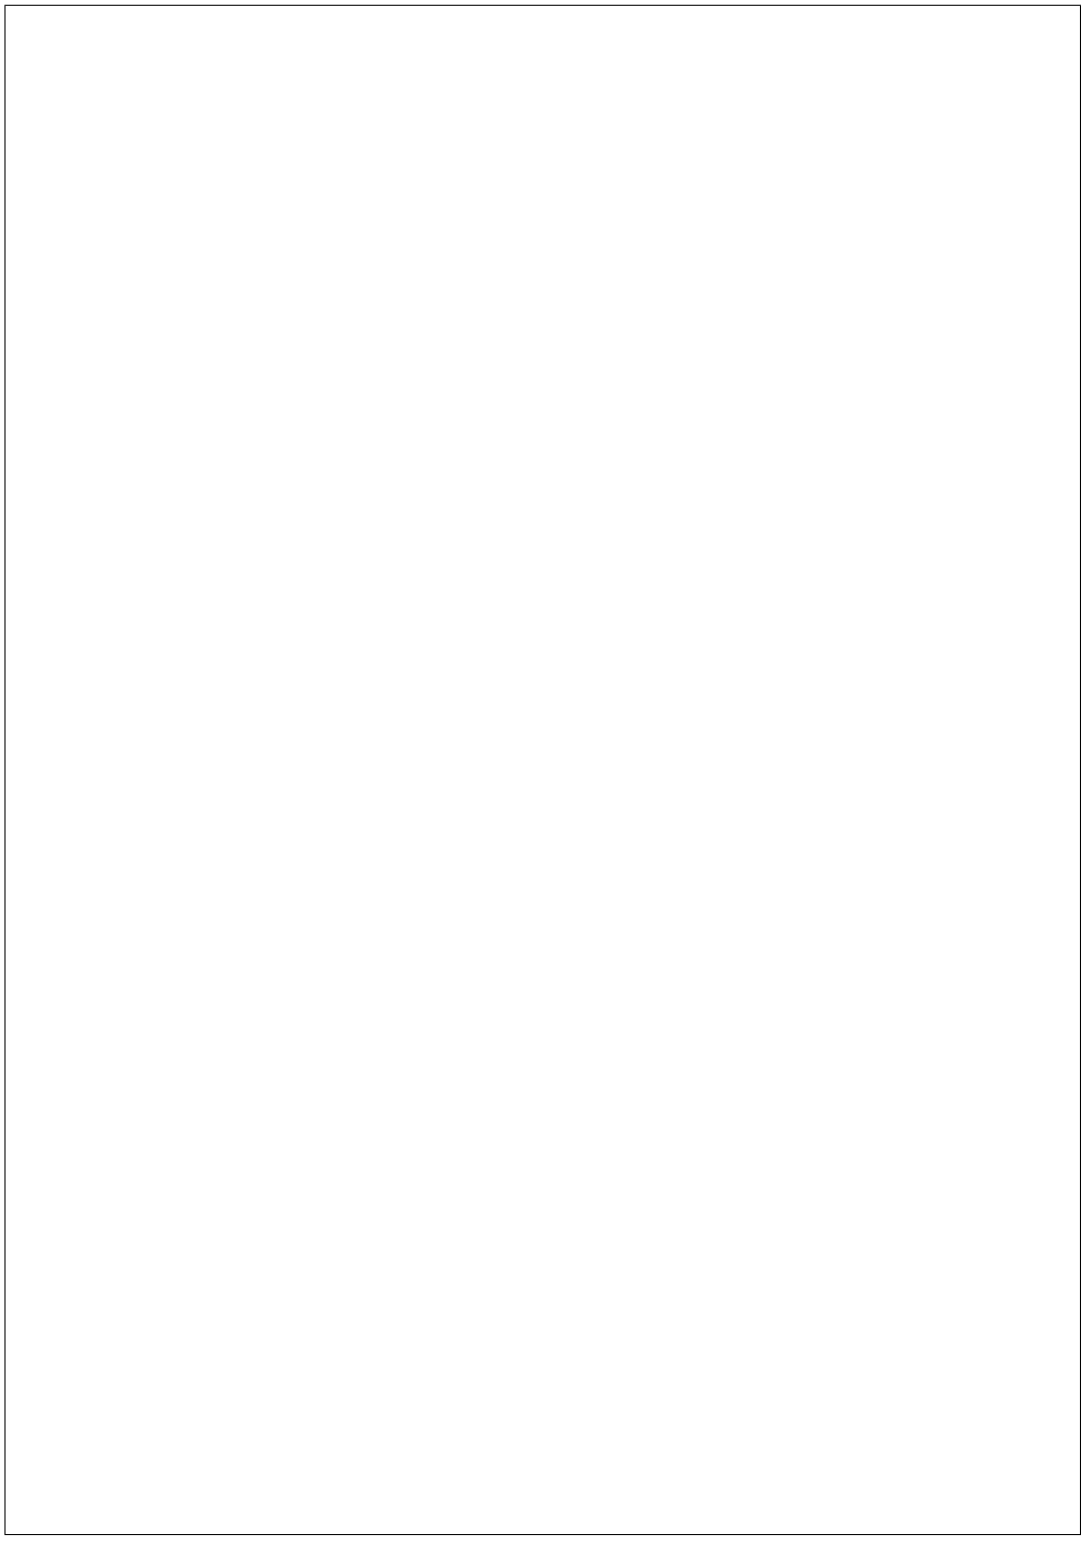

22.4.2 For the following graphs of vector fields, determine whether the divergence is positive, negative or zero.

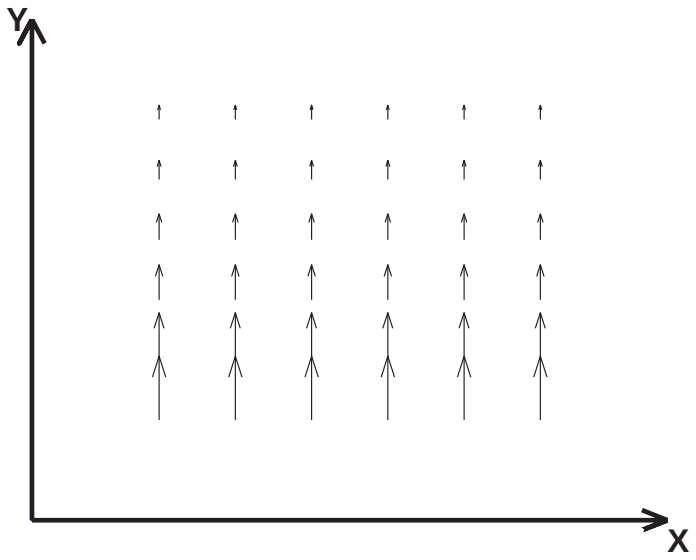

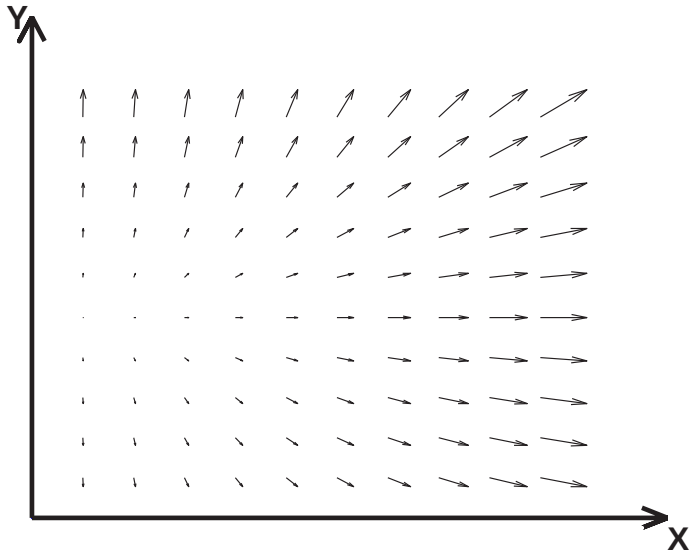

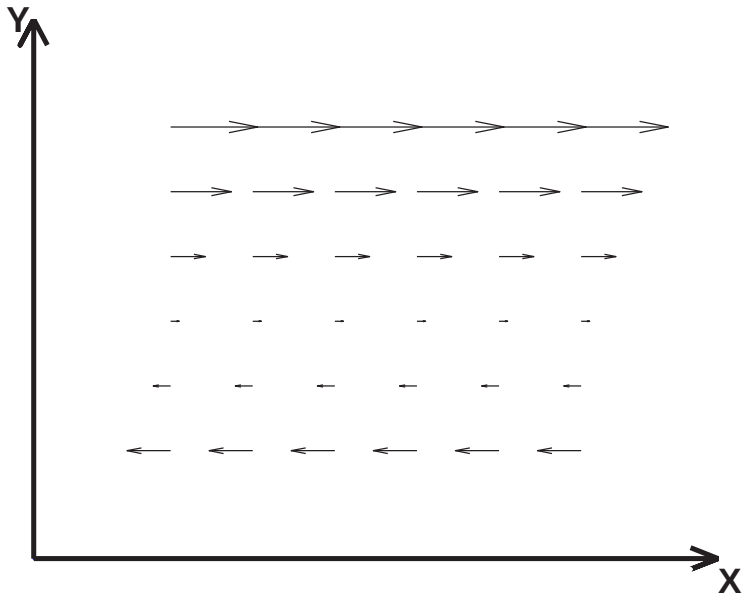
Y
X

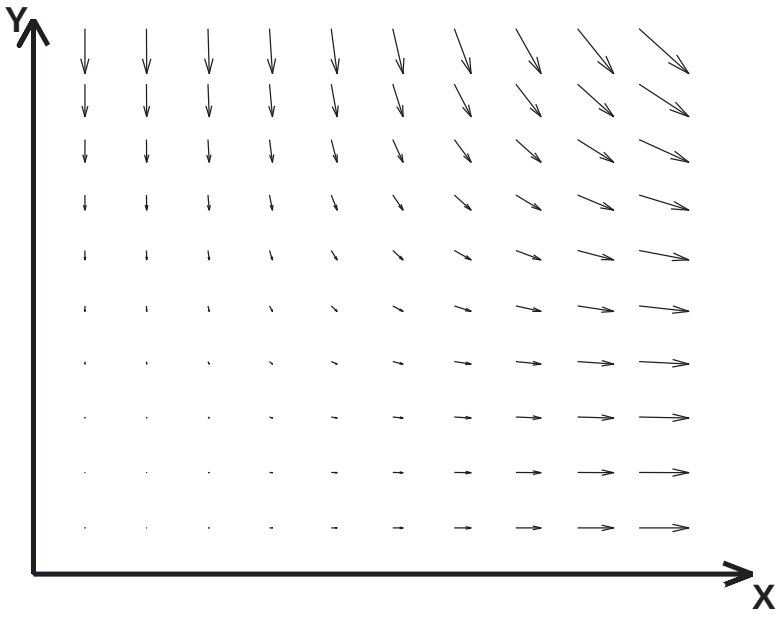
Y
X

23 Parametrisation of surfaces in $\mathbb{R}^3$

By the end of this section, you should be able to answer the following questions:

- What does it mean to parametrise a surface in $\mathbb{R}^3$?
- How do you parametrise certain surfaces?

23.1 Parametric surfaces

We have already seen two ways of representing a surface in $\mathbb{R}^3$: explicitly as $z = f(x, y)$ or implicitly as $F(x, y, z) = 0$.

Another way of representing a surface S in $\mathbb{R}^3$ is by a parametrisation. This is where the coordinate variables are functions of two parameters u and v:

$$x = x(u, v), \quad y = y(u, v), \quad z = z(u, v)$$

and the vector

$$\boldsymbol{r}(u, v) = x(u, v)\boldsymbol{i} + y(u, v)\boldsymbol{j} + z(u, v)\boldsymbol{k}$$

traces out the surface as u, v vary over some region D in the "u-v plane". So for every point (u, v) in D, there corresponds a point on the surface S.

The following diagram shows the point P on the surface S which corresponds to the point (u, v) in the region D in the u-v plane. As (u, v) moves around all points in D, the point P moves around in S, tracing out the entire surface.

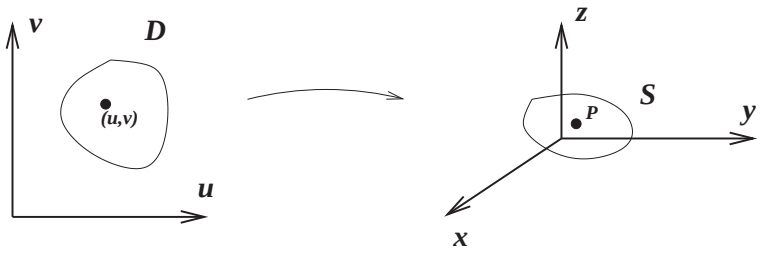

Note that a surface defined explicitly by $z = f(x, y)$ is equivalent to a parametrisation

$$\boldsymbol{r}(x, y) = x\boldsymbol{i} + y\boldsymbol{j} + f(x, y)\boldsymbol{k},$$

where we treat the coordinate variables x and y as the parameters. Note that we have not specified any bounds on the variables. Often the challenge is to not only find suitable functions for a parametrisation, but for a finite surface to determine bounds on the parameters.

23.2 Parametrising surfaces using cylindrical and spherical coordinates

We can use our knowledge of cylindrical and spherical coordinates to parametrise certain surfaces with which these coordinates are naturally associated.

Recall cylindrical coordinates:

$$x = r\cos\theta, \quad y = r\sin\theta, \quad z = z.$$

Setting exactly one of the cylindrical coordinates to a constant value necessarily gives a parametric surface.

Setting $z = 2$ with $0 \leq \theta \leq 2\pi$, $0 \leq r \leq 3$ describes a disc of radius 3, centred at the z axis lying in the plane $z = 2$:

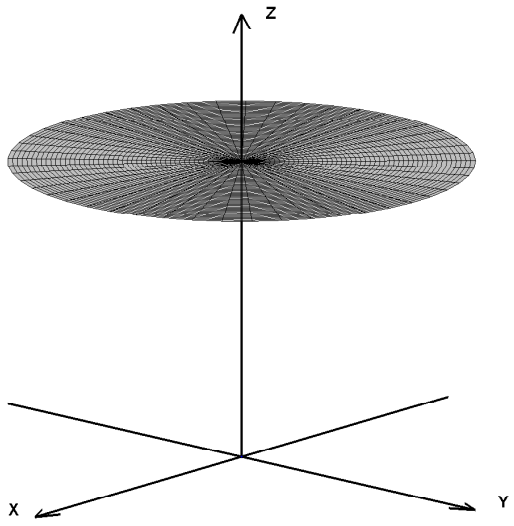

Setting $r = 5$ with $0 \leq \theta \leq 2\pi$, $1 \leq z \leq 3$ describes the surface of a cylinder of radius 5 and of height 2 between $z = 1$ and $z = 3$:

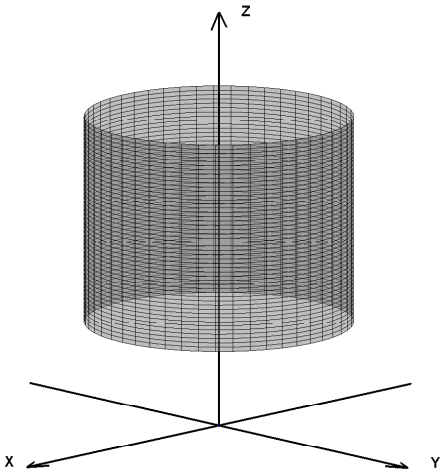

Setting $\theta = \pi/2$ with $2 \leq z \leq 4$, $0 \leq r \leq 1$ describes a rectangle lying in the y-z plane. Another description of the same surface would be $x = 0,\ \{(y, z) \mid 0 \leq y \leq 1,\ 2 \leq z \leq 4\}$:

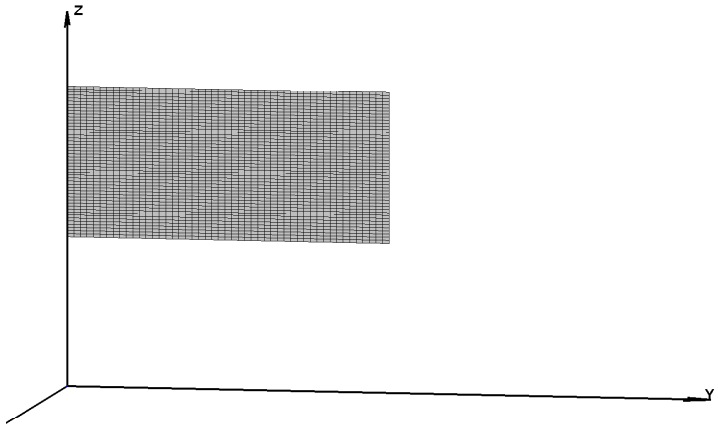

23.2.1 Parametrise the paraboloid $z = 1 - x^2 - y^2$ for $z \geq 0$.

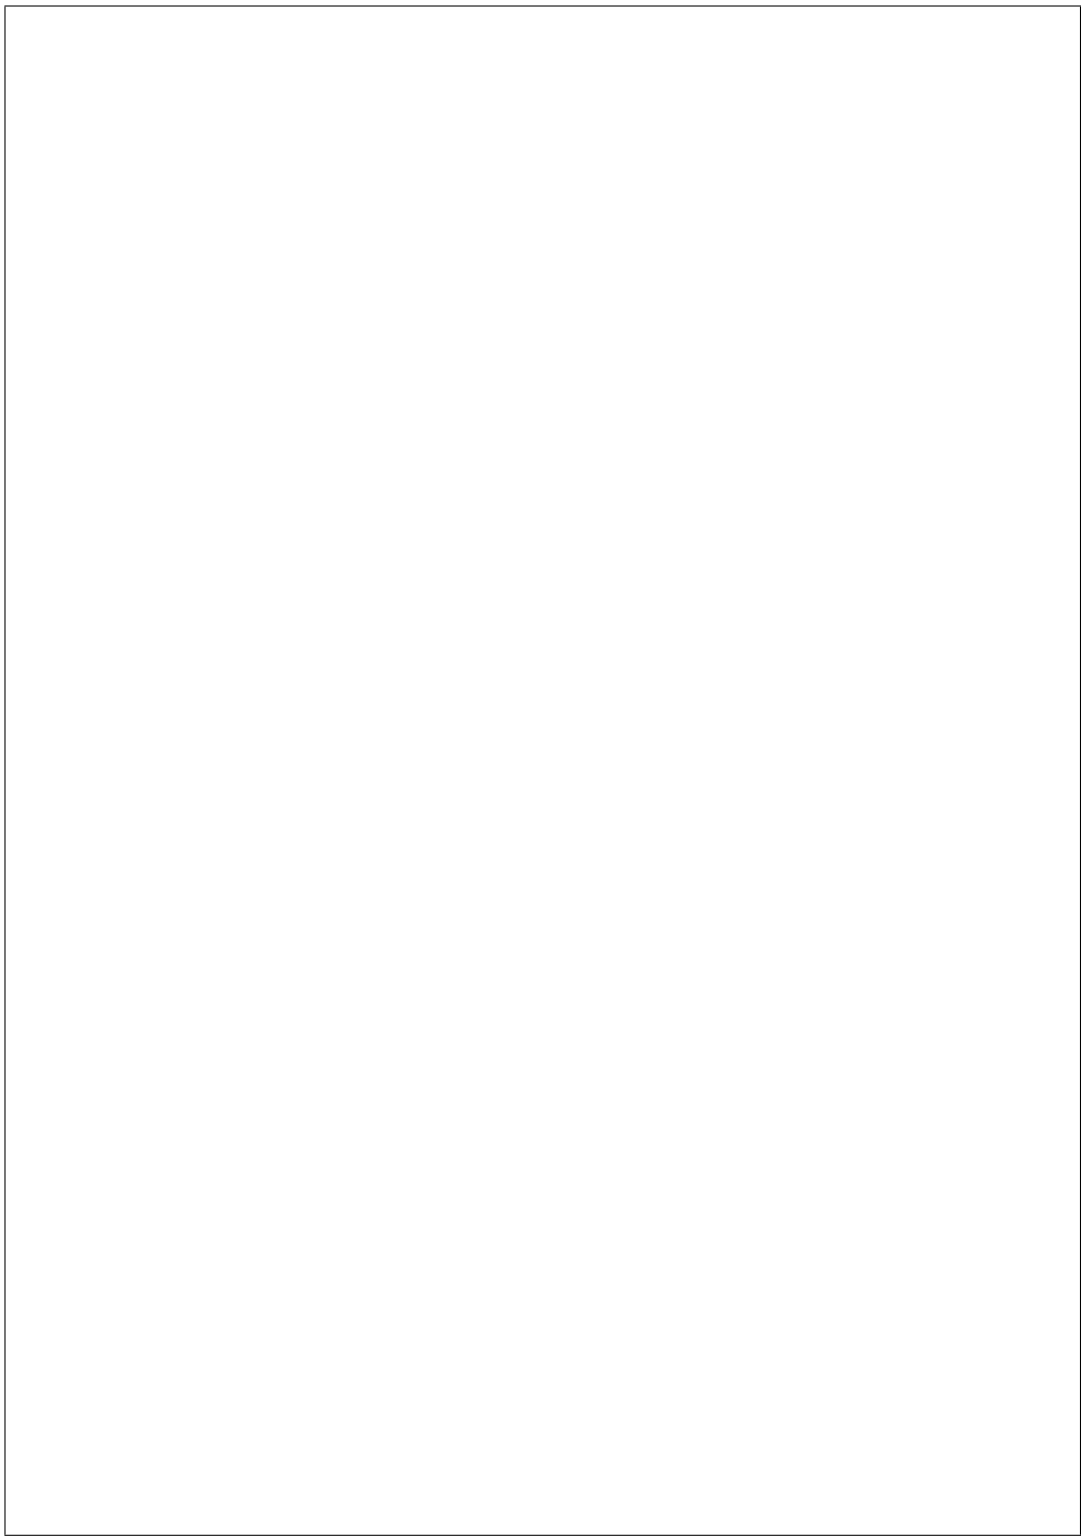

Recall spherical coordinates: $x = r\cos\theta\sin\phi, \quad y = r\sin\theta\sin\phi, \quad z = r\cos\phi.$

Setting exactly one of the spherical coordinates to a constant value necessarily gives a parametric surface.

Setting $r = 2$ with $0 \leq \theta \leq 2\pi$, $0 \leq \phi \leq \pi$ describes the surface of a sphere of radius 2 centred at the origin:

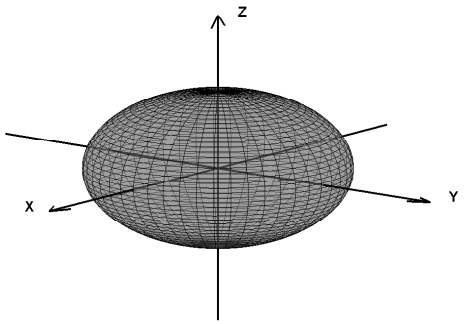

Setting $\phi = \pi/3$ with $0 \leq r \leq 2$, $0 \leq \theta \leq 2\pi$ describes the open cone with angle $\pi/3$ to the positive z-axis, the "mouth" of which lies on the sphere of radius 2 and with vertex located at the origin:

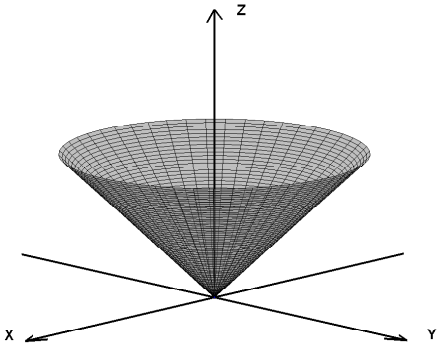

Setting $\theta = 0$ with $0 \leq r \leq 3$, $0 \leq \phi \leq \pi$ describes the half disc of radius 3 lying in the x-z plane:

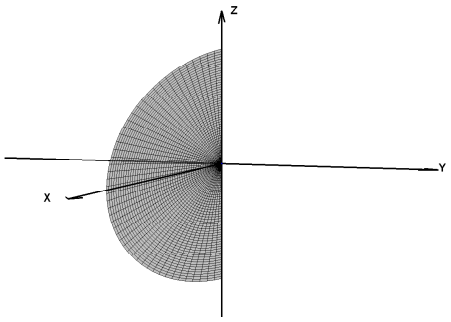

23.2.2 Parametrise the part of the sphere $x^2 + y^2 + z^2 = 16$ that lies between the planes $z = 2$ and $z = -2$.

23.3 Tangent planes

Let S be a surface parametrised by

$$\boldsymbol{r}(u,v) = x(u,v)\boldsymbol{i} + y(u,v)\boldsymbol{j} + z(u,v)\boldsymbol{k}.$$

Here we find the tangent plane to S at a point P specified by $\boldsymbol{r}(a,b)$.

There are two important families of curves on S. One where u is a constant, the other where v is a constant. The diagram below shows the relationship between horizontal and vertical lines in D (in the u-v plane) and curves on S.

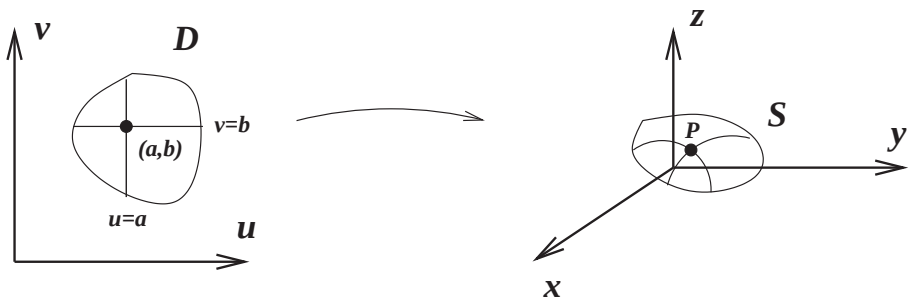

Setting $u = a$ defines a curve on S parametrised by $\boldsymbol{r}(a,v)$, for all values of v such that (a,v) lies in D. A tangent vector to this curve at P is

$$\boldsymbol{r}_v = \frac{\partial x}{\partial v}(a,b)\boldsymbol{i} + \frac{\partial y}{\partial v}(a,b)\boldsymbol{j} + \frac{\partial z}{\partial v}(a,b)\boldsymbol{k}.$$

Similarly setting $v = b$ defines another curve on S parametrised by $\boldsymbol{r}(u,b)$. A tangent vector to this curve at P is

$$\boldsymbol{r}_u = \frac{\partial x}{\partial u}(a,b)\boldsymbol{i} + \frac{\partial y}{\partial u}(a,b)\boldsymbol{j} + \frac{\partial z}{\partial u}(a,b)\boldsymbol{k}.$$

If $\boldsymbol{r}_u$ and $\boldsymbol{r}_v$ are continuous and $\boldsymbol{r}_u \times \boldsymbol{r}_v$ is never $\boldsymbol{0}$ inside D (we make an exception for points on the boundary of D), we call the surface *smooth* (it has no "kinks").

For a smooth surface, $\boldsymbol{r}_u \times \boldsymbol{r}_v$ is a normal vector at any point inside D. This vector evaluated at $(u,v) = (a,b)$ is also normal to the tangent plane at the point $P = (x(a,b), y(a,b), z(a,b))$.

The equation of the tangent plane at P is given by

$$(\boldsymbol{r}_u(a,b) \times \boldsymbol{r}_v(a,b)) \cdot ((x\boldsymbol{i} + y\boldsymbol{j} + z\boldsymbol{k}) - \boldsymbol{r}(a,b)) = 0.$$

23.3.1 Find the tangent plane to the surface parametrised by $\boldsymbol{r}(u,v) = u^2\boldsymbol{i} + v^2\boldsymbol{j} + (u+2v)\boldsymbol{k}$ at the point $(1,1,3)$.

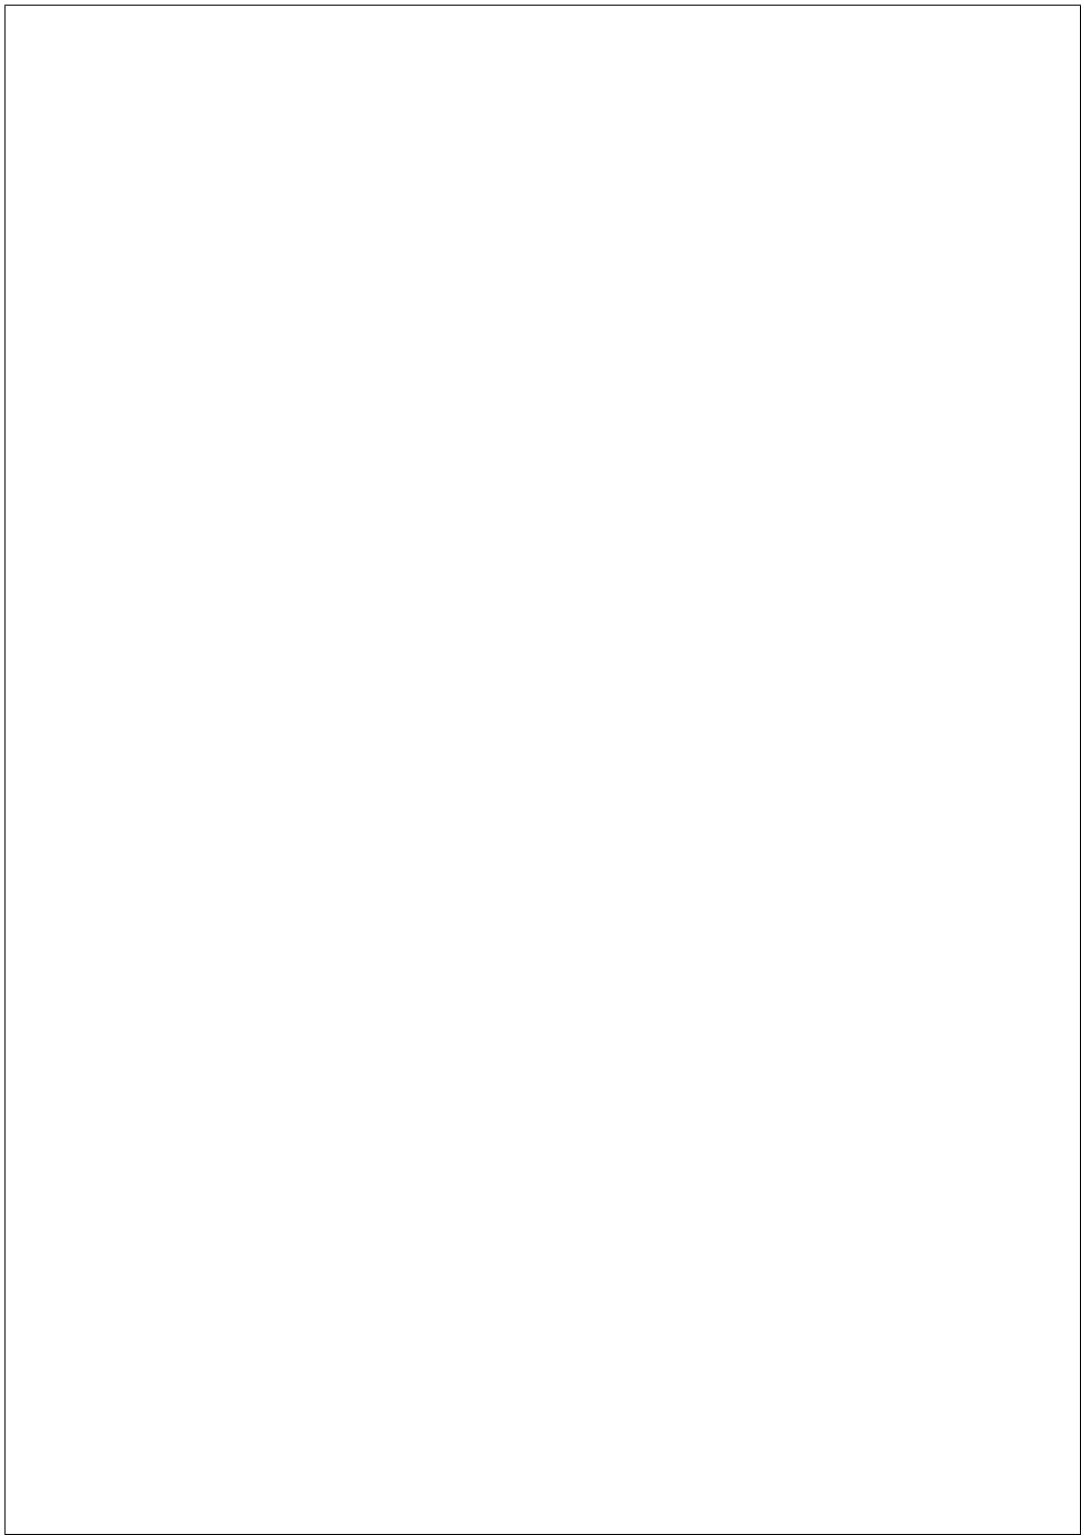

24 Surface integrals

By the end of this section, you should be able to answer the following questions:

- What is a surface integral?
- How do you calculate the area of a parametric surface?
- How do you use surface integrals in applications such as calculating the mass of a "surface lamina" and finding the average temperature over a surface.

24.1 Area of a parametric surface

Let S be a smooth parametric surface given by

$$\boldsymbol{r}(u,v) = x(u,v)\boldsymbol{i} + y(u,v)\boldsymbol{j} + z(u,v)\boldsymbol{k},$$

where we assume for simplicity that the parameter domain is a rectangle in the u-v plane. To calculate the area of S, we work through the following steps:

1. Partition S into small patches.
2. Approximate each patch by a parallelogram lying in the tangent plane to the corner of the patch closest to the u-v origin.
3. Calculate the area ΔS of each parallelogram and add them to give an approximation to the area of S.
4. Take the limit as the dimensions of $\Delta S \to 0$ to obtain an exact expression for the area.

Let's have a closer look at each step.

1. A partition of S into patches will correspond to a partition of D (in the u-v plane) into small rectangles.

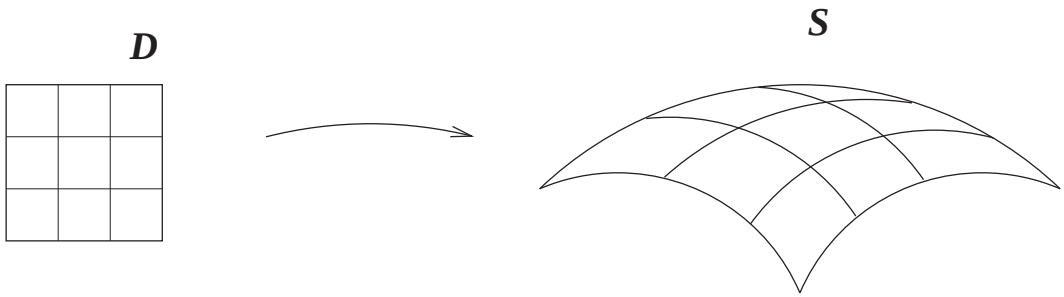

   The dimensions of the rectangles in D will be $\Delta u \Delta v$.

2. Let one of the edges of a single patch be defined from parameter values (u, v) to $(u + \Delta u, v)$.

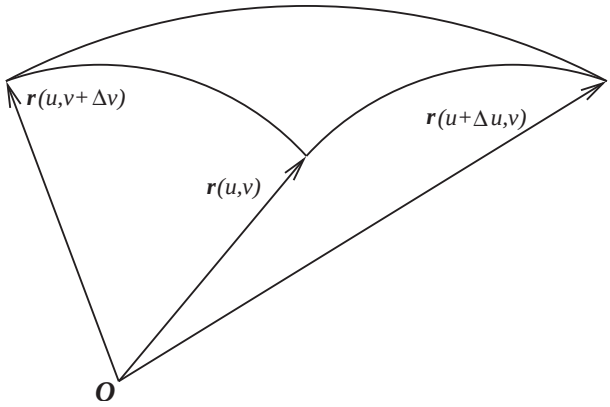

   Using Pythagoras' law in three dimensions, we can approximate the length of this edge as

$$\begin{aligned} \text{length} &\approx \sqrt{(\Delta x)^2 + (\Delta y)^2 + (\Delta z)^2} \\ &= \left(\sqrt{\left(\frac{\Delta x}{\Delta u}\right)^2 + \left(\frac{\Delta y}{\Delta u}\right)^2 + \left(\frac{\Delta z}{\Delta u}\right)^2} \right) \Delta u \\ &\approx |\boldsymbol{r}_u| \Delta u, \end{aligned}$$

   where in this case we have used $\Delta x = x(u+\Delta u, v) - x(u, v)$ etc (ie. the change is only in u). Similarly, for an edge of patch running from parameter values (u, v) to $(u, v + \Delta v)$ the length of that edge will be approximately $|\boldsymbol{r}_v| \Delta v$.

At the corner of the patch corresponding to parameter values (u, v), we can define the two vectors $\boldsymbol{r}_u \Delta u$ and $\boldsymbol{r}_v \Delta v$ which form two sides of a parallelogram, the side lengths of which coincide with our approximations to the lengths of the edges of the patch.

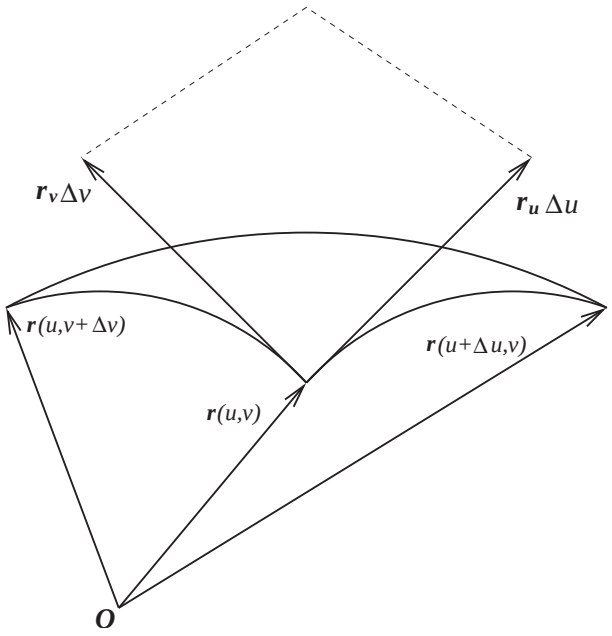

3. The vector $(\boldsymbol{r}_u\ \Delta u) \times (\boldsymbol{r}_v\ \Delta v)$ is normal to the surface (and hence the tangent plane) at that point. Its magnitude gives the area of the parallelogram we use to approximate the area of the patch ΔS. We then have

$$\Delta S \approx |\boldsymbol{r}_u \times \boldsymbol{r}_v|\ \Delta u\ \Delta v.$$

Adding these approximations for each patch in S gives us an approximation to the area of S:

$$\text{area of } S \approx \sum_i \Delta S_i = \sum_i |\boldsymbol{r}_{u_i} \times \boldsymbol{r}_{v_i}|\ \Delta u_i\ \Delta v_i.$$

4. Finally taking the limit as $\Delta u, \Delta v \to 0$ we obtain

$$\text{surface area } = \iint_S dS = \iint_D |\boldsymbol{r}_u \times \boldsymbol{r}_v|\ du\ dv.$$

24.1.1 Application: find the surface area of the paraboloid $z = 1 - x^2 - y^2$ for $z \geq 0$.

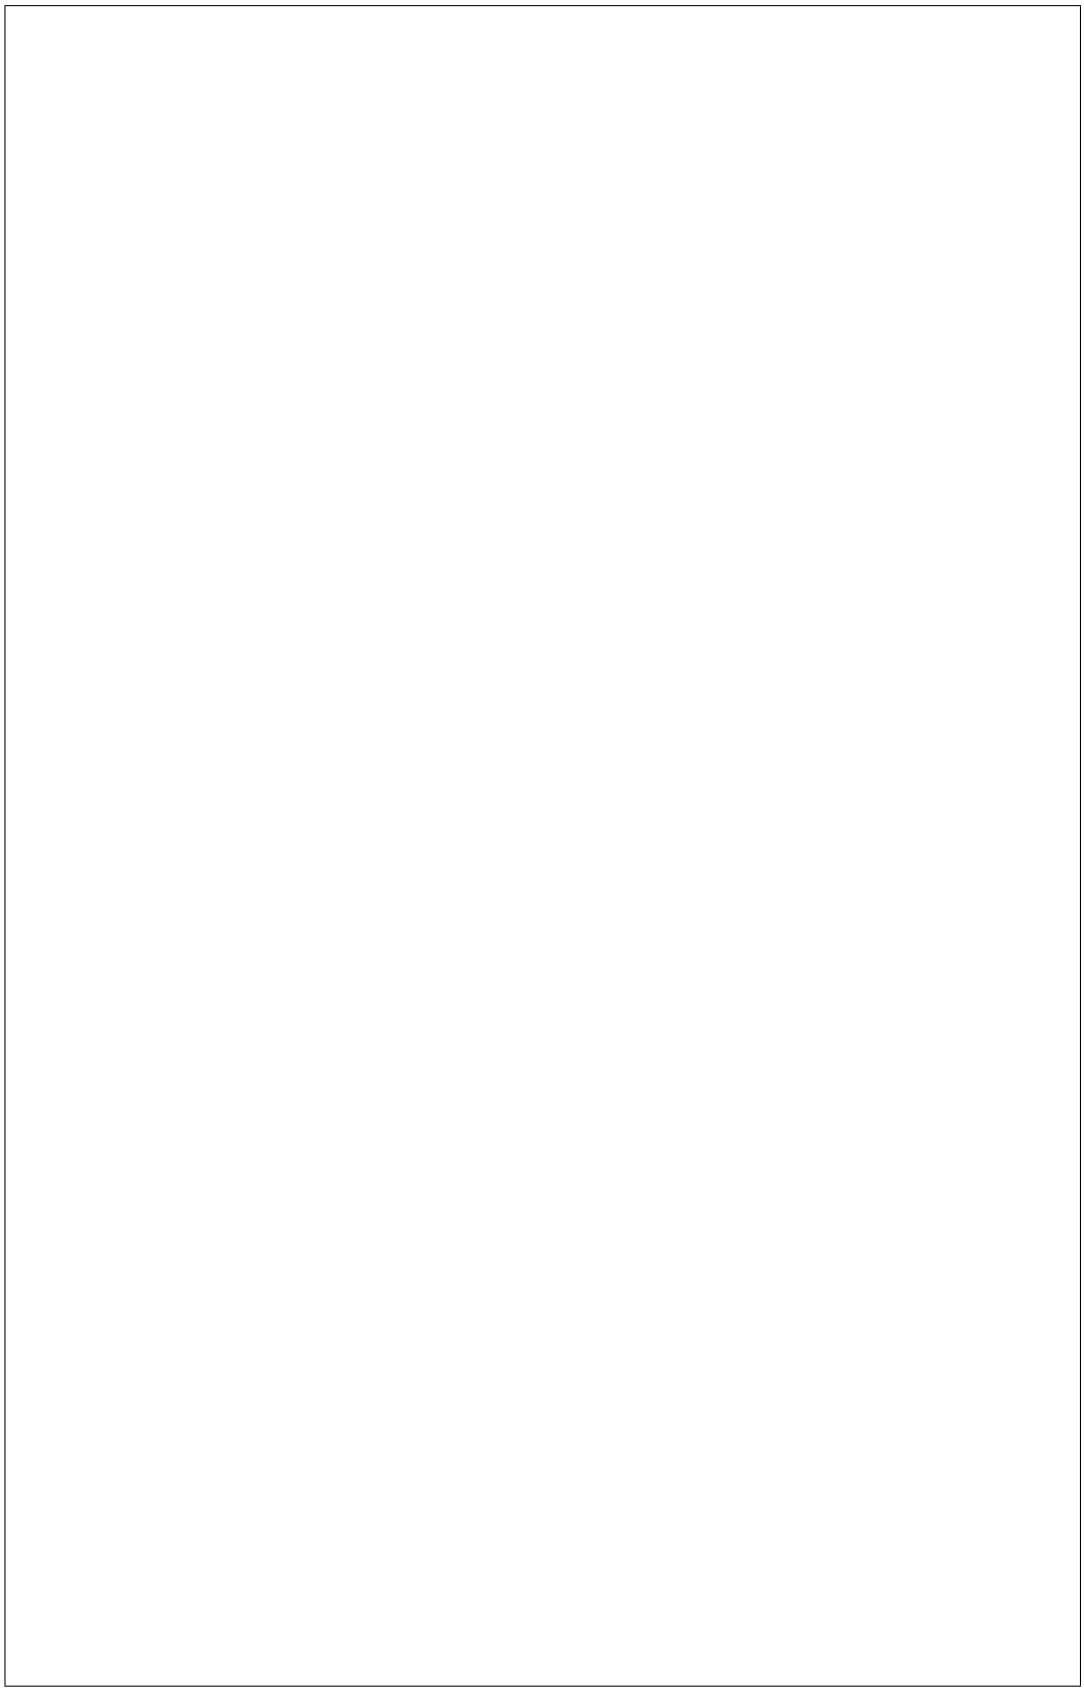

24.2 More on calculating surface integrals, applications

Let $f(x, y, z)$ be a scalar function in $\mathbb{R}^3$. We can define the surface integral of f over a smooth parametric surface S in $\mathbb{R}^3$ as

$$\iint_S f(x, y, z)\ dS = \iint_D f(\boldsymbol{r}(u, v))|\boldsymbol{r}_u \times \boldsymbol{r}_v|\ du\ dv.$$

Surface integrals and double integrals have similar applications. Indeed, a double integral is merely a special case of a surface integral where the surface lies entirely in the x-y plane.

For example, if a thin sheet has the shape of a surface S and the mass density at the point (x, y, z) is $\rho(x, y, z)$, then the mass of the sheet is given by a surface integral:

$$\text{mass of sheet } = \iint_S \rho(x, y, z)\ dS.$$

Another application is in calculating the average value of a function over a surface. Let S be a smooth surface in $\mathbb{R}^3$. Then the average value of the function $f(x, y, z)$ over that surface is given by

$$\text{average value over surface } = \frac{1}{\text{area of } S} \iint_S f(x, y, z)\ dS.$$

If the surface S is a closed surface, it is convention to write

$$\oiint_S f(x, y, z)\ dS$$

to represent the surface integral.

If S is a finite union of smooth surfaces S_1, S_2,..., S_n that intersect only at their boundaries, then

$$\iint_S f(x, y, z)\ dS = \iint_{S_1} f(x, y, z)\ dS + \iint_{S_2} f(x, y, z)\ dS + \ldots + \iint_{S_n} f(x, y, z)\ dS.$$

Closed surfaces are often unions of smooth surfaces as demonstrated in the following example.

24.2.1 **The function $T(x, y, z) = x^2 + y^2 + z^2 + 4$ gives the temperature at any point (x, y, z) on the surface of a solid hemisphere of radius 1 centred at the origin, defined for $z \geq 0$. Find the average temperature over the surface.**

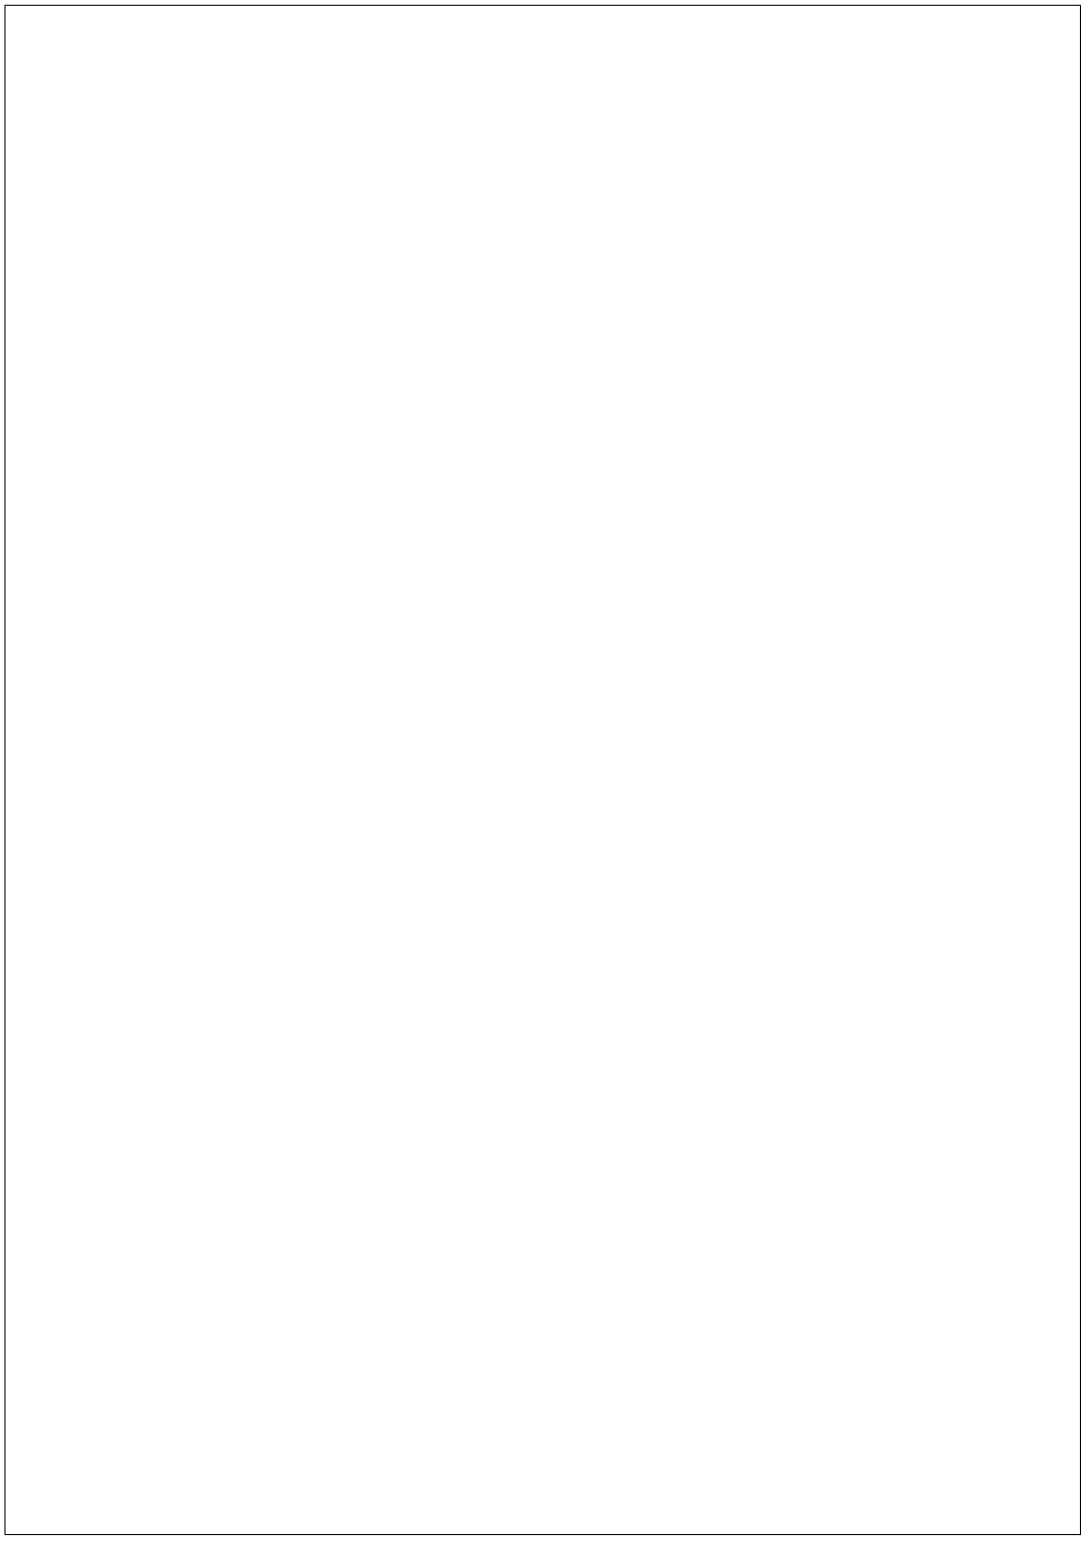

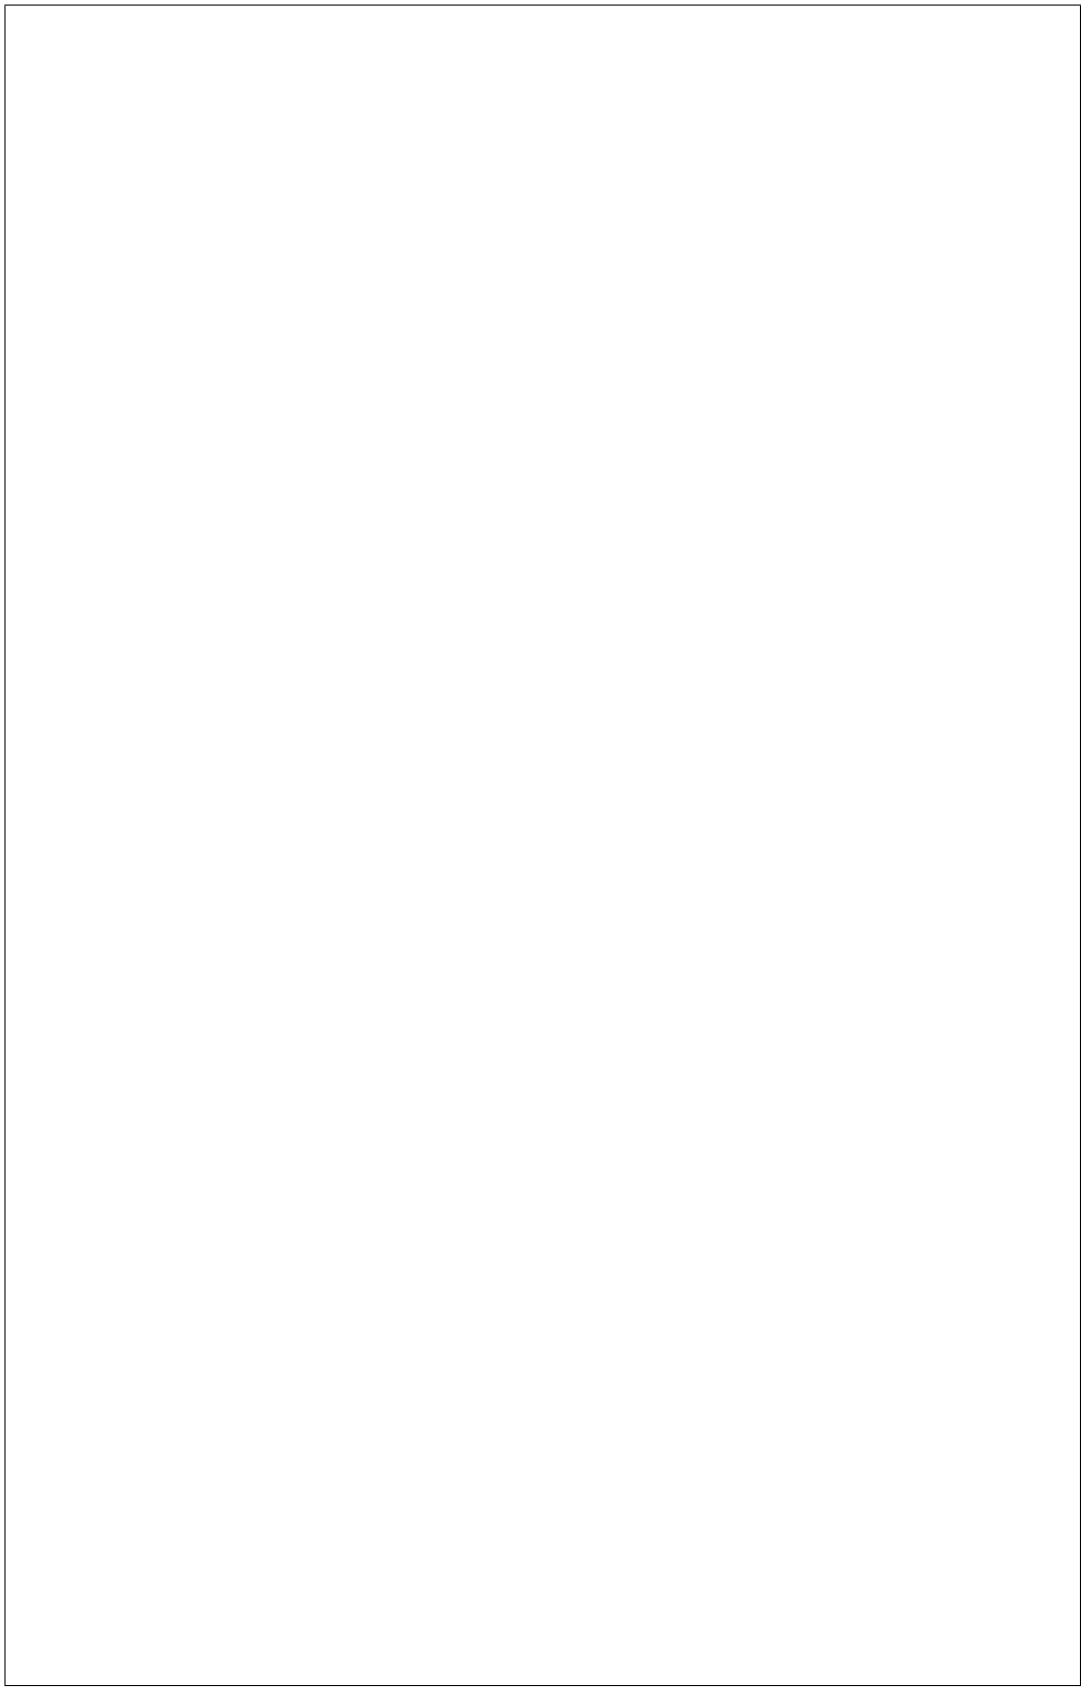

25 Flux integrals and Gauss' divergence theorem

By the end of this section, you should be able to answer the following questions:

- What is a flux integral?
- How do you use a flux integral to calculate the flux of a vector field across a surface?
- What is Gauss' divergence theorem and under what conditions can it be applied?
- How do you apply Gauss' divergence theorem?

We have already been introduced to the idea of flux of a variable vector field across a curve (in $\mathbb{R}^2$) and the flux of a constant vector field across rectangular surfaces (in $\mathbb{R}^3$). In this section we look at calculating the flux across smoothly parametric surfaces.

25.1 Orientable surfaces

Let S be a smooth surface. If we can choose a unit vector that is normal to S at every point so that $\boldsymbol{n}$ varies continuously over S, we call S an *orientable* surface. The choice of $\boldsymbol{n}$ provides S with an *orientation*. There are only ever two possible orientations.

An example of an orientable surface is the surface of a sphere. The two possible orientations are out of the sphere or into the sphere.

An example of a non-orientable surface is a Möbius strip (see Stewart page 1139).

The orientation of a surface is important when considering flux through that surface. The orientation we choose is always the direction of positive flux.

25.2 The flux integral

For a vector field $\boldsymbol{v}(x, y, z)$, we are interested in the flux of $\boldsymbol{v}$ across a smooth orientable parametric surface S in $\mathbb{R}^3$, parametrised by $\boldsymbol{r}(u, v)$, with u and v defined over some domain D. Let $\boldsymbol{n}(u, v)$ be a unit vector normal to the surface S which defines the orientation of the surface (and hence the direction of positive flux).

It would be most convenient to consider the context of fluid flow with $\boldsymbol{v}(x, y, z)$ being the velocity of a fluid at the point (x, y, z).

To calculate the flux through S, we work through the following steps:

1. Partition S into small patches.
2. Approximate each patch by a parallelogram lying in the tangent plane to the corner of the patch closest to the u-v origin.
3. Approximate the flux through each parallelogram of approximate area ΔS and add them to give an approximation to the total flux through S.
4. Take the limit as the dimensions of $\Delta S \to 0$ to obtain an exact expression for the flux.

Let's have a closer look at these steps.

1,2. Steps 1 and 2 are exactly the same as steps 1 and 2 on page 156-157 of our calculation of surface area.

3. We approximate the flux through one patch by treating $\boldsymbol{v}$ as constant over the patch (ie. the patch is small enough for this to be a decent approximation). Since we have already approximated the shape of the patch as a parallelogram, we need to work out the flux of a constant vector through a parallelogram.

   To this end, consider the parallelogram defined by the two (non-parallel) vectors $\boldsymbol{a}$ and $\boldsymbol{b}$. If we take the area of the patch to be ΔS, it can be seen from the diagram below that the flux (volume per unit time if $\boldsymbol{v}$ is velocity) passing through the parallelogram is

$$\text{flux across parallelogram } \approx \boldsymbol{v} \cdot \boldsymbol{n}\, \Delta S.$$

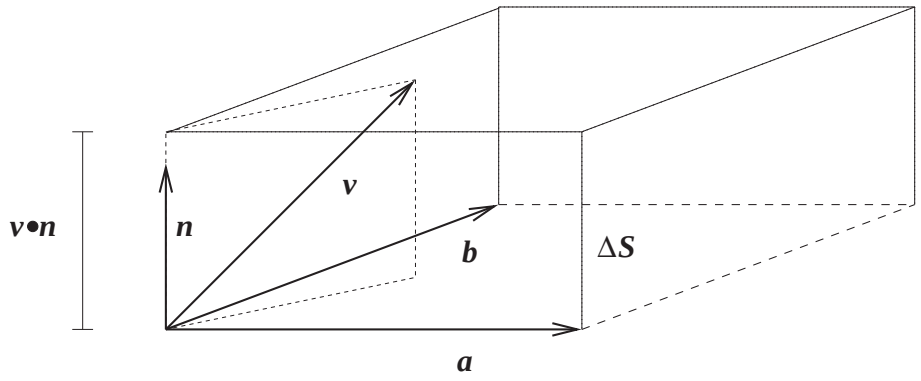

   We can take $\boldsymbol{n} = \dfrac{\boldsymbol{a} \times \boldsymbol{b}}{|\boldsymbol{a} \times \boldsymbol{b}|}$, and the area of the parallelogram is $|\boldsymbol{a} \times \boldsymbol{b}| \approx \Delta S$. We then have

$$\text{flux across parallelogram } \approx \boldsymbol{v} \cdot \frac{\boldsymbol{a} \times \boldsymbol{b}}{|\boldsymbol{a} \times \boldsymbol{b}|}\, |\boldsymbol{a} \times \boldsymbol{b}| = \boldsymbol{v} \cdot (\boldsymbol{a} \times \boldsymbol{b}).$$

As shown previously, a patch of surface can be approximated by a parallelogram determined by the two vectors $\boldsymbol{r}_u \Delta u$ and $\boldsymbol{r}_v \Delta v$. Hence we have

$$\text{flux across one patch } \approx \boldsymbol{v} \cdot \boldsymbol{n}\ \Delta S = \boldsymbol{v} \cdot (\boldsymbol{r}_u \times \boldsymbol{r}_v) \Delta u \Delta v.$$

Note that we need to check that the vector $\boldsymbol{r}_u \times \boldsymbol{r}_v$ points in the direction of positive flux. If not, we use $\boldsymbol{r}_v \times \boldsymbol{r}_u$.

Adding these approximations over the entire surface S, we obtain

$$\text{flux across } S \approx \sum_i \boldsymbol{v}_i \cdot \boldsymbol{n}_i\ \Delta S_i = \sum_i \boldsymbol{v}(u_i, v_i) \cdot (\boldsymbol{r}_{u_i} \times \boldsymbol{r}_{v_i}) \Delta u_i \Delta v_i.$$

4. To obtain an exact expression for the flux across S we take the limit as $\Delta u, \Delta v \to 0$.

$$\text{flux across } S = \iint_S \boldsymbol{v} \cdot \boldsymbol{n}\ dS = \iint_D \boldsymbol{v} \cdot (\boldsymbol{r}_u \times \boldsymbol{r}_v)\ du\ dv.$$

This expression is called a *flux integral* and is used to calculate the flux of any vector field across a smooth orientable surface, not just fluids with a given velocity field.

25.2.1 Calculate the net outward flux of $\boldsymbol{F}(x, y, z) = z\boldsymbol{i} + y\boldsymbol{j} + x\boldsymbol{k}$ across the surface of the cylindrical solid given by $\{(x, y, z) \mid x^2 + y^2 \leq 1,\ 0 \leq z \leq 2\}$.

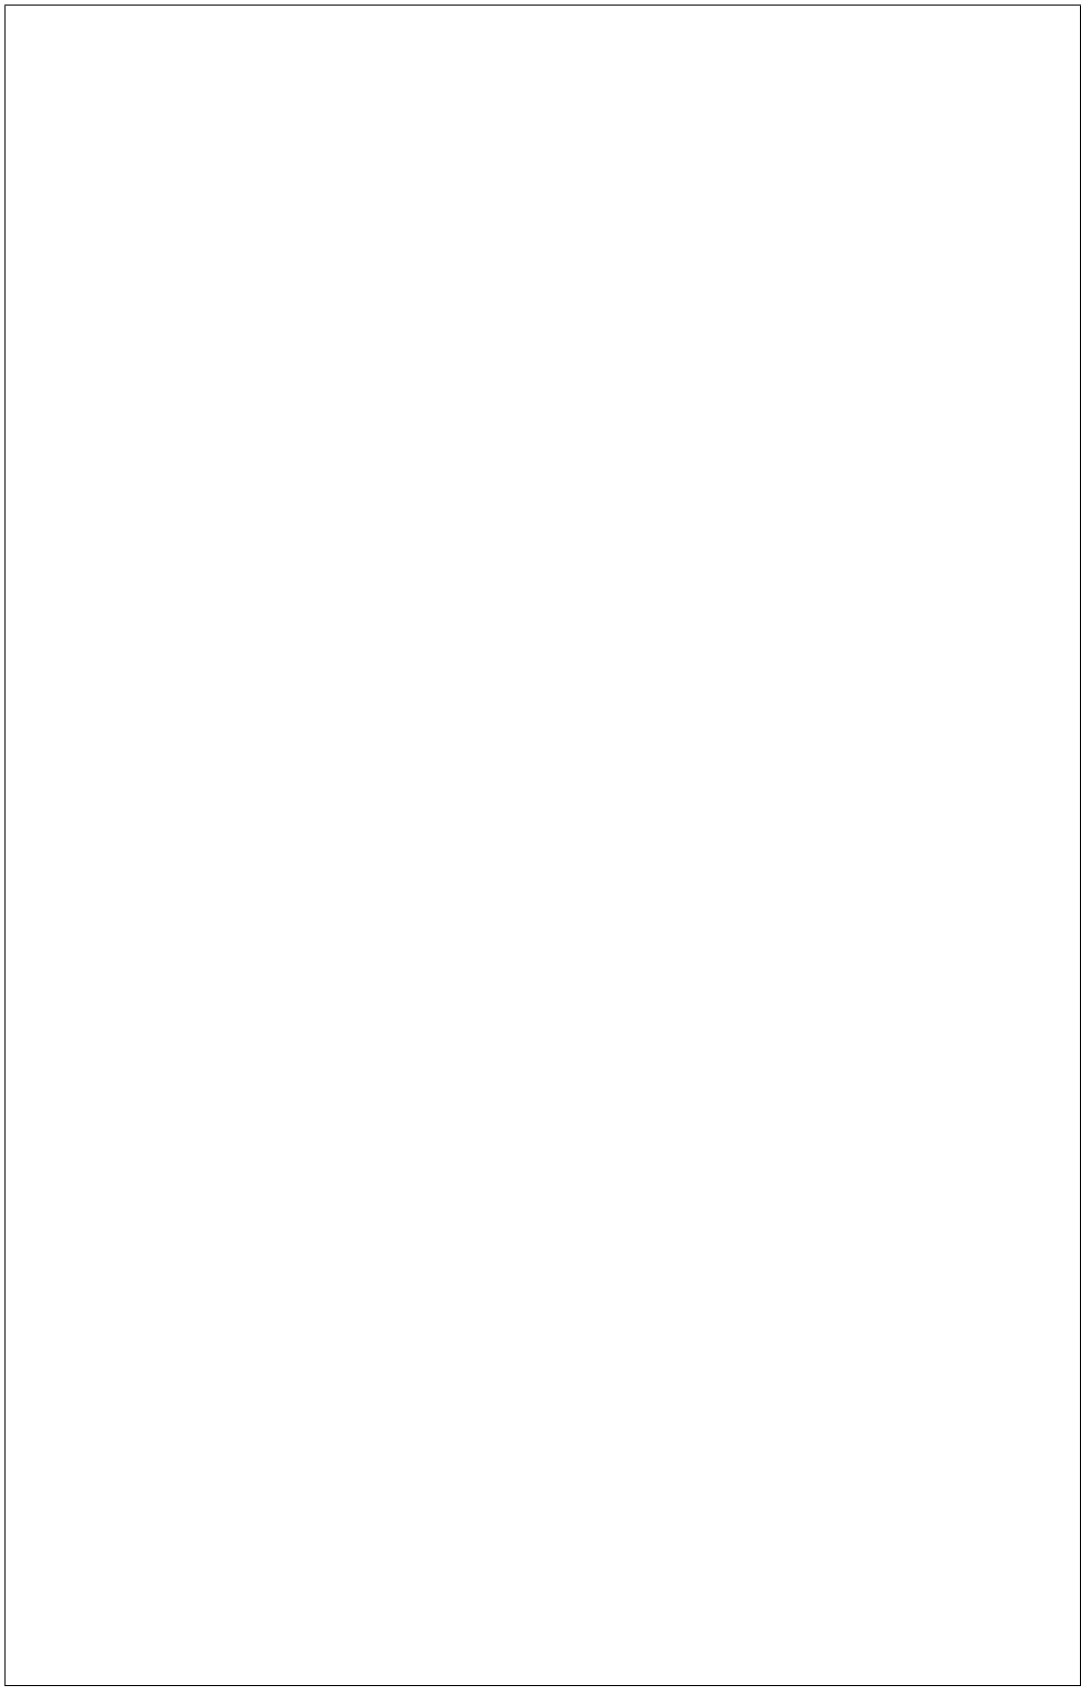

25.3 Gauss' divergence theorem

On page 142 we saw the flux form of Green's theorem:

$$\oint_C \boldsymbol{v}(x,y)\cdot\boldsymbol{n}\;dS = \iint_D \operatorname{div}(\boldsymbol{v}(x,y))\;dA.$$

The left hand side is essentially a flux integral in two dimensions, with $\boldsymbol{n}$ being an outwardly pointing unit normal vector to the curve C. The right hand side was derived from our realisation of the divergence as the "flux density".

It would be natural to ask if it is possible to extend this result to three dimensions.

Given a vector field in three dimensions, $\boldsymbol{F}(x,y,z)$, we have seen that the net outward flux across a closed, smooth, orientable surface S is given by $\oiint_S \boldsymbol{F}\cdot\boldsymbol{n}\;dS$, where $\boldsymbol{n}$ is an outwardly pointing unit normal.

We have also seen that its divergence ($\operatorname{div}\boldsymbol{F}$) can be viewed as the flux density, so

$$\operatorname{div}\boldsymbol{F} = \lim_{\Delta V\to 0}\frac{\text{flux of }\boldsymbol{F}\text{ out of }\Delta V}{\Delta V}.$$

Hence we expect to be able to calculate the net outward flux across a closed, smooth, orientable surface S as the triple integral of the flux density (ie. $\operatorname{div}\boldsymbol{F}$) over the volume enclosed by S.

Indeed, this is true, with $\boldsymbol{F}$ and S subject to certain conditions. The result is known as *Gauss' divergence theorem*:

Let S be a piecewise smooth, orientable, closed surface enclosing a region V in $\mathbb{R}^3$. Let $\boldsymbol{F}(x,y,z)$ be a vector field whose component functions are continuous and have continuous partial derivatives in V. Then

$$\oiint_S \boldsymbol{F}\cdot\boldsymbol{n}\;dS = \iiint_V \operatorname{div}(\boldsymbol{F})dV,$$

where $\boldsymbol{n}$ is the outwardly directed unit normal to S.

This theorem connects the flux of a vector field out of a volume with the flux through its surface. It says that we can calculate the net outward flux either as a closed surface integral, or as a triple integral.

25.3.1 Use Gauss' divergence theorem to calculate the net outward flux of $\boldsymbol{F}(x, y, z) = z\boldsymbol{i} + y\boldsymbol{j} + x\boldsymbol{k}$ across the surface of the cylindrical solid given by $\{(x, y, z) \mid x^2 + y^2 \leq 1,\ 0 \leq z \leq 2\}$.

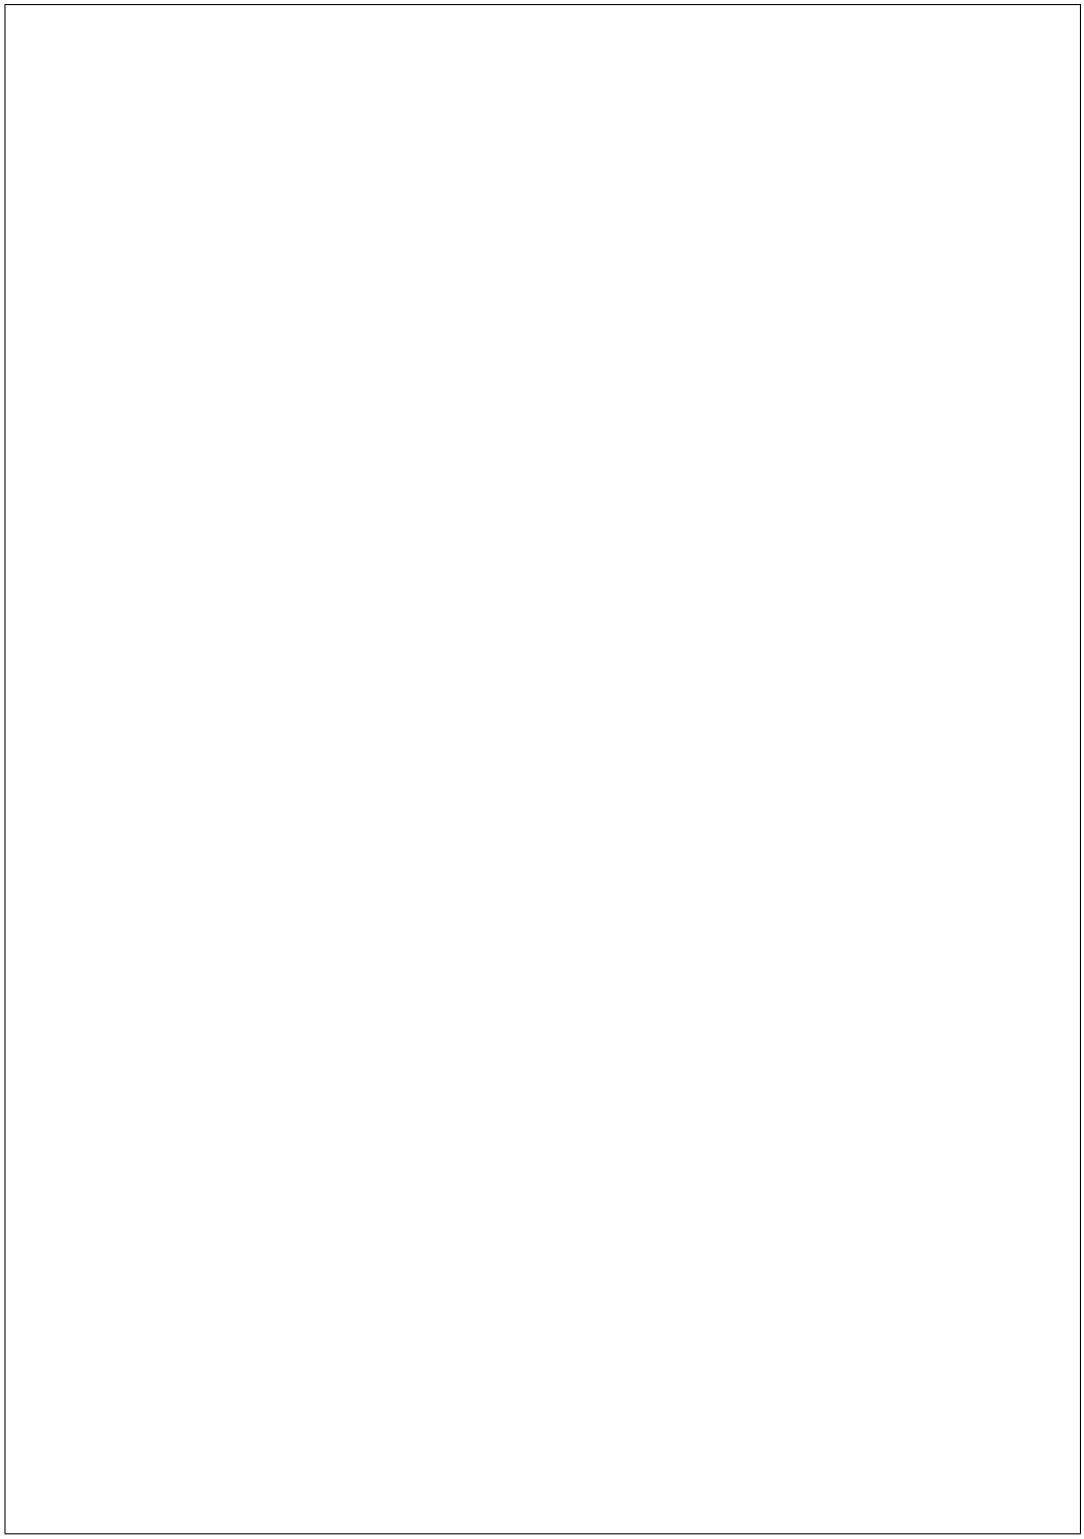

25.3.2 Application: net outward flux of an electric field across any closed surface enclosing the origin.

The electric field (force per unit charge) of a charge Q located at the origin is given by

$$\boldsymbol{E}(\boldsymbol{r}) = \frac{1}{4\pi\varepsilon_0}\frac{Q}{|\boldsymbol{r}|^2}\left(\frac{\boldsymbol{r}}{|\boldsymbol{r}|}\right) = \frac{Q}{4\pi\varepsilon_0}\frac{\boldsymbol{r}}{|\boldsymbol{r}|^3},$$

where ε_0 is a physical constant (permittivity), and as usual $\boldsymbol{r} = x\boldsymbol{i}+y\boldsymbol{j}+z\boldsymbol{k}$. We are interested in calculating the net outward flux across any closed surface S enclosing the origin.

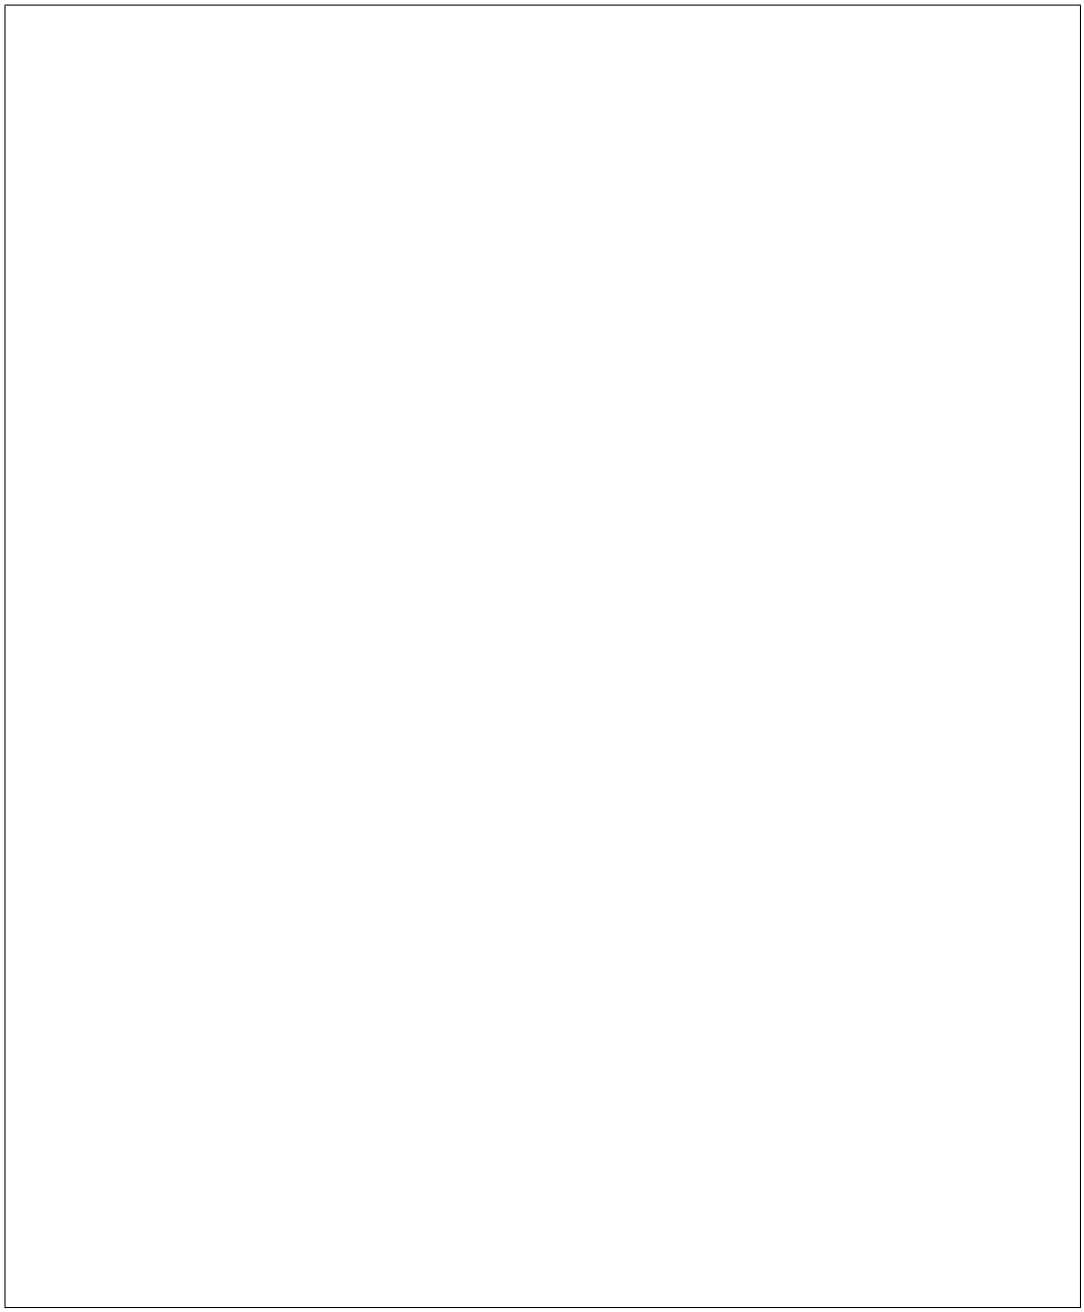

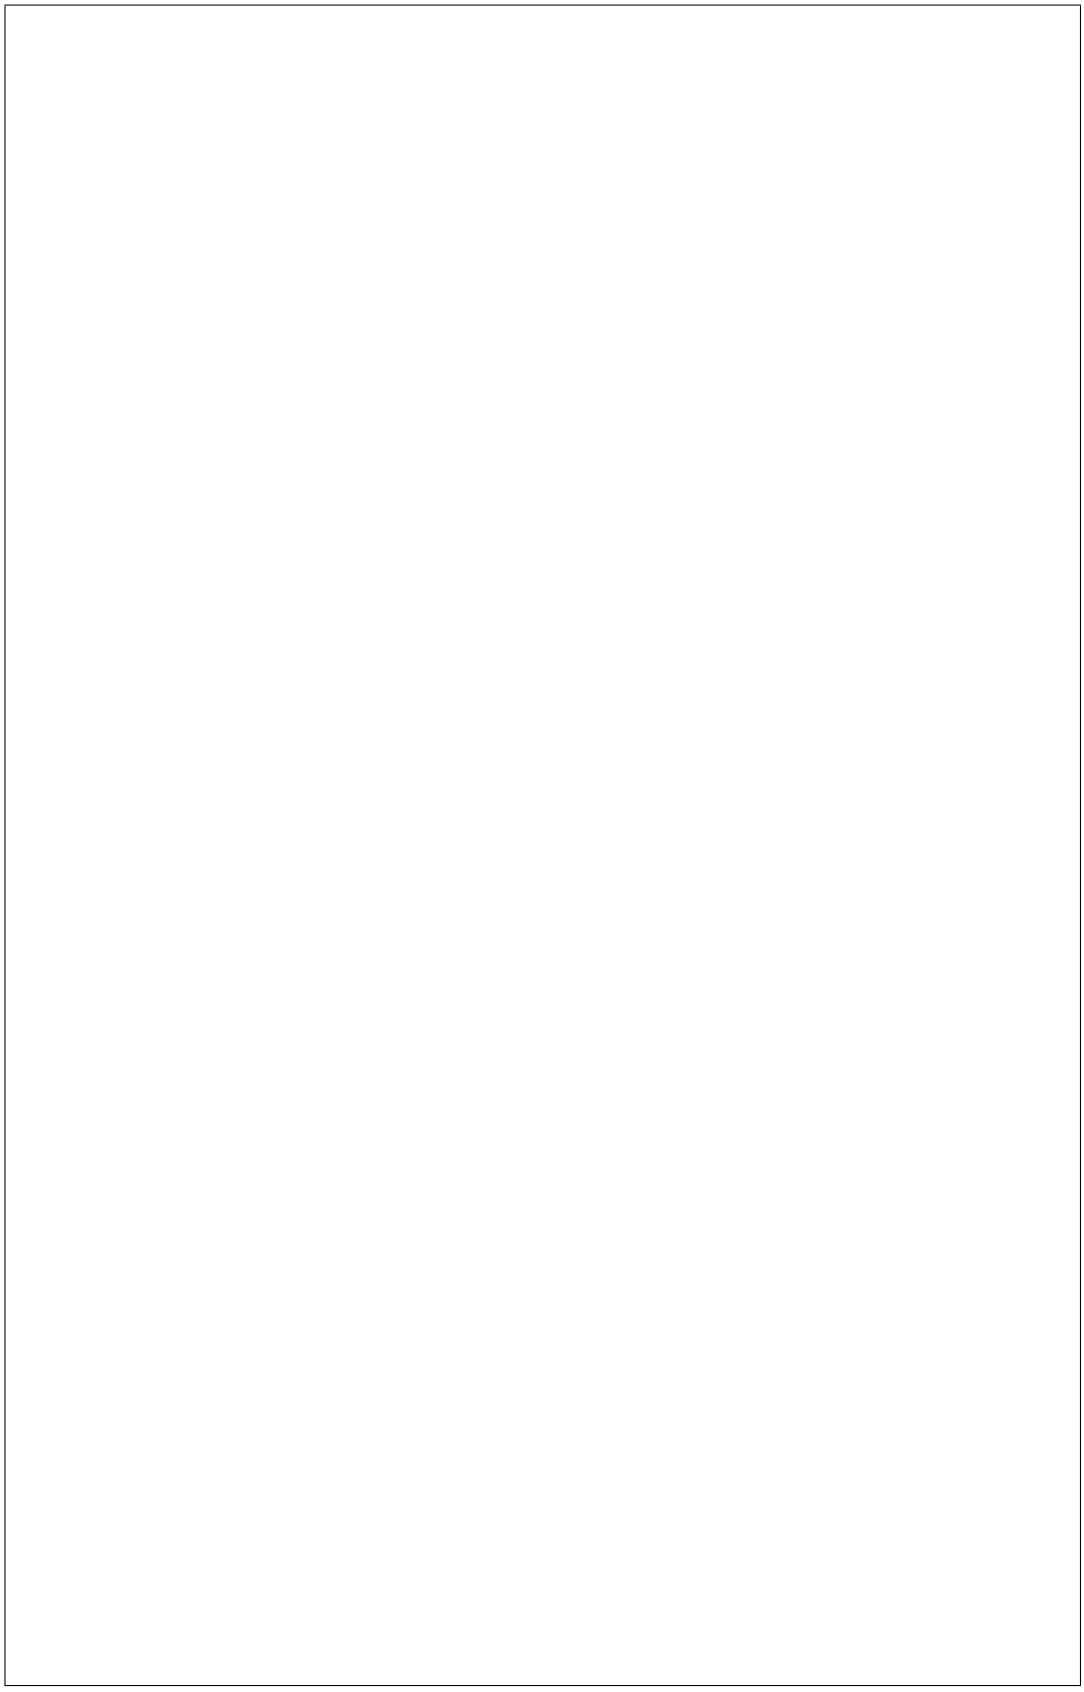

26 Curl of a vector field

By the end of this section, you should be able to answer the following questions:

- How do you calculate the curl of a given vector field?
- What is the significance of curl?
- How do you test whether or not a given three dimensional vector field is conservative?

26.1 Calculating curl

If (x, y, z) is a right handed Cartesian coordinate system and $\boldsymbol{v}(x, y, z) = v_1\boldsymbol{i} + v_2\boldsymbol{j} + v_3\boldsymbol{k}$ is a differentiable vector field, then define

$$\text{curl}(\boldsymbol{v}) = \nabla \times \boldsymbol{v} = \begin{vmatrix} \boldsymbol{i} & \boldsymbol{j} & \boldsymbol{k} \\ \frac{\partial}{\partial x} & \frac{\partial}{\partial y} & \frac{\partial}{\partial z} \\ v_1 & v_2 & v_3 \end{vmatrix}$$

$$= \left(\frac{\partial v_3}{\partial y} - \frac{\partial v_2}{\partial z}\right)\boldsymbol{i} + \left(\frac{\partial v_1}{\partial z} - \frac{\partial v_3}{\partial x}\right)\boldsymbol{j} + \left(\frac{\partial v_2}{\partial x} - \frac{\partial v_1}{\partial y}\right)\boldsymbol{k}.$$

Note that $\text{curl}(\boldsymbol{v})$ is a vector field.

26.1.1 Example: let $\boldsymbol{v} = yz^2\boldsymbol{i} + zx^2\boldsymbol{j} + xy^2\boldsymbol{k}$. Find curl$(\boldsymbol{v})$.

26.2 Understanding curl

For the rotation of a rigid body about a fixed axis with angular velocity $\boldsymbol{w}$, the velocity at a point P, whose position vector is $\boldsymbol{r}$, is given by $\boldsymbol{v} = \boldsymbol{w} \times \boldsymbol{r}$.

If we choose the axis of rotation to be the z-axis, then $\boldsymbol{w} = \omega \boldsymbol{k}$. Calculate curl($\boldsymbol{v}$).

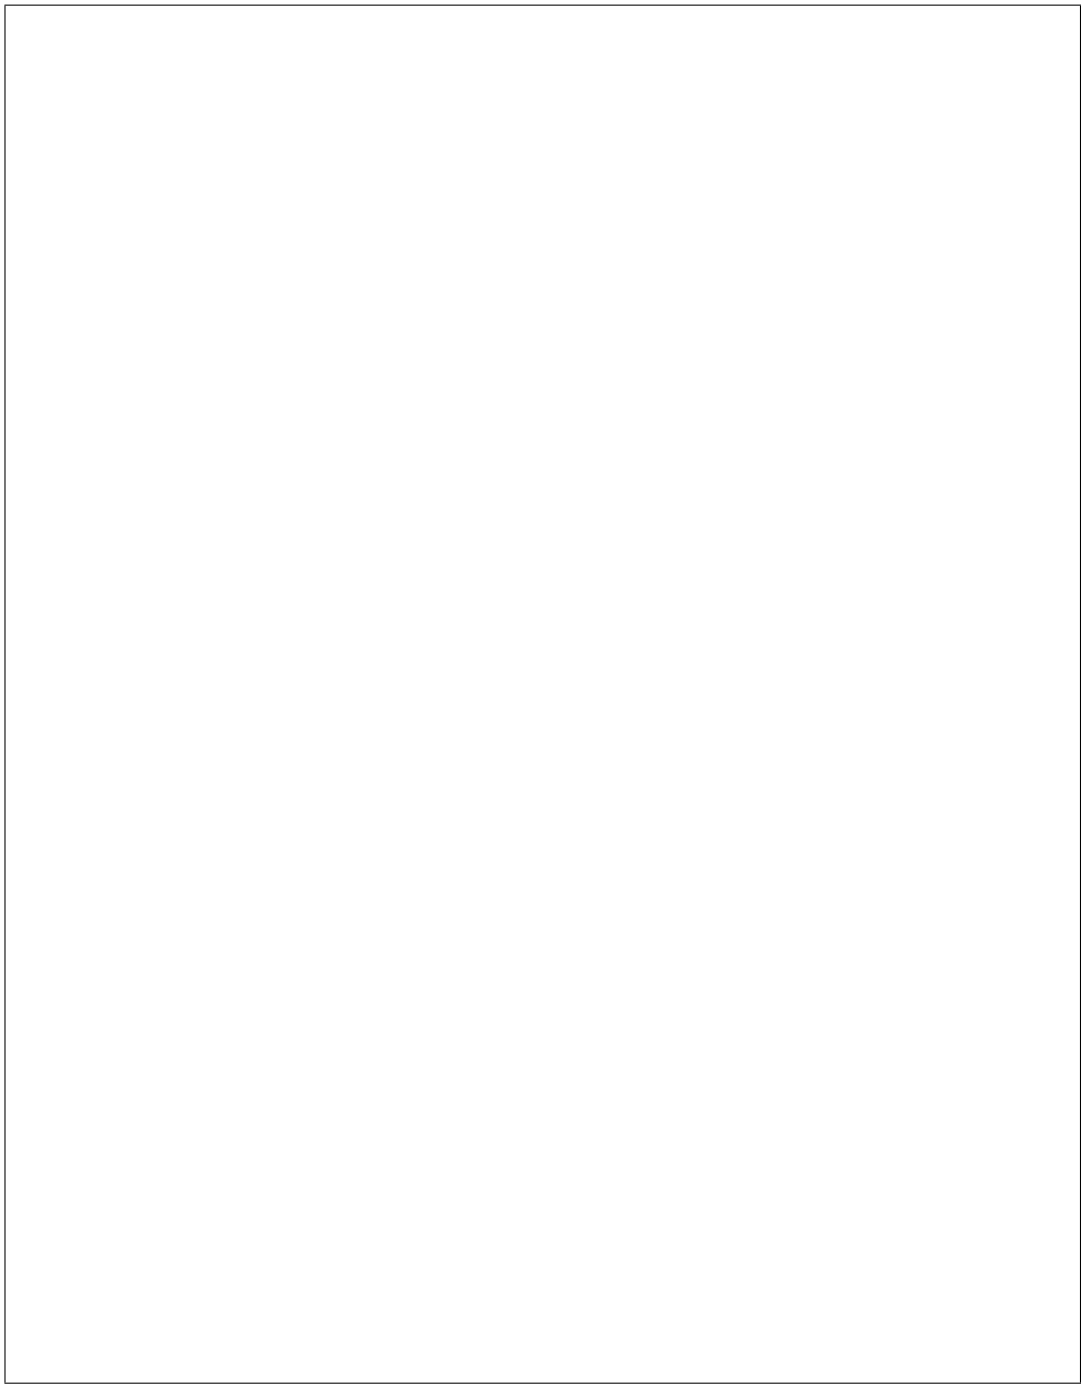

In general, curl($\boldsymbol{v}$) characterises the rotation of a vector field. We will investigate this further in the next section.

26.3 Conservative fields revisited

It turns out that the curl of a vector field is exactly what we need to generalise the result at the bottom of page 126 to three dimensions.

Show that if $\boldsymbol{F}$ is a conservative vector field, then $\text{curl}\boldsymbol{F} = \mathbf{0}$.

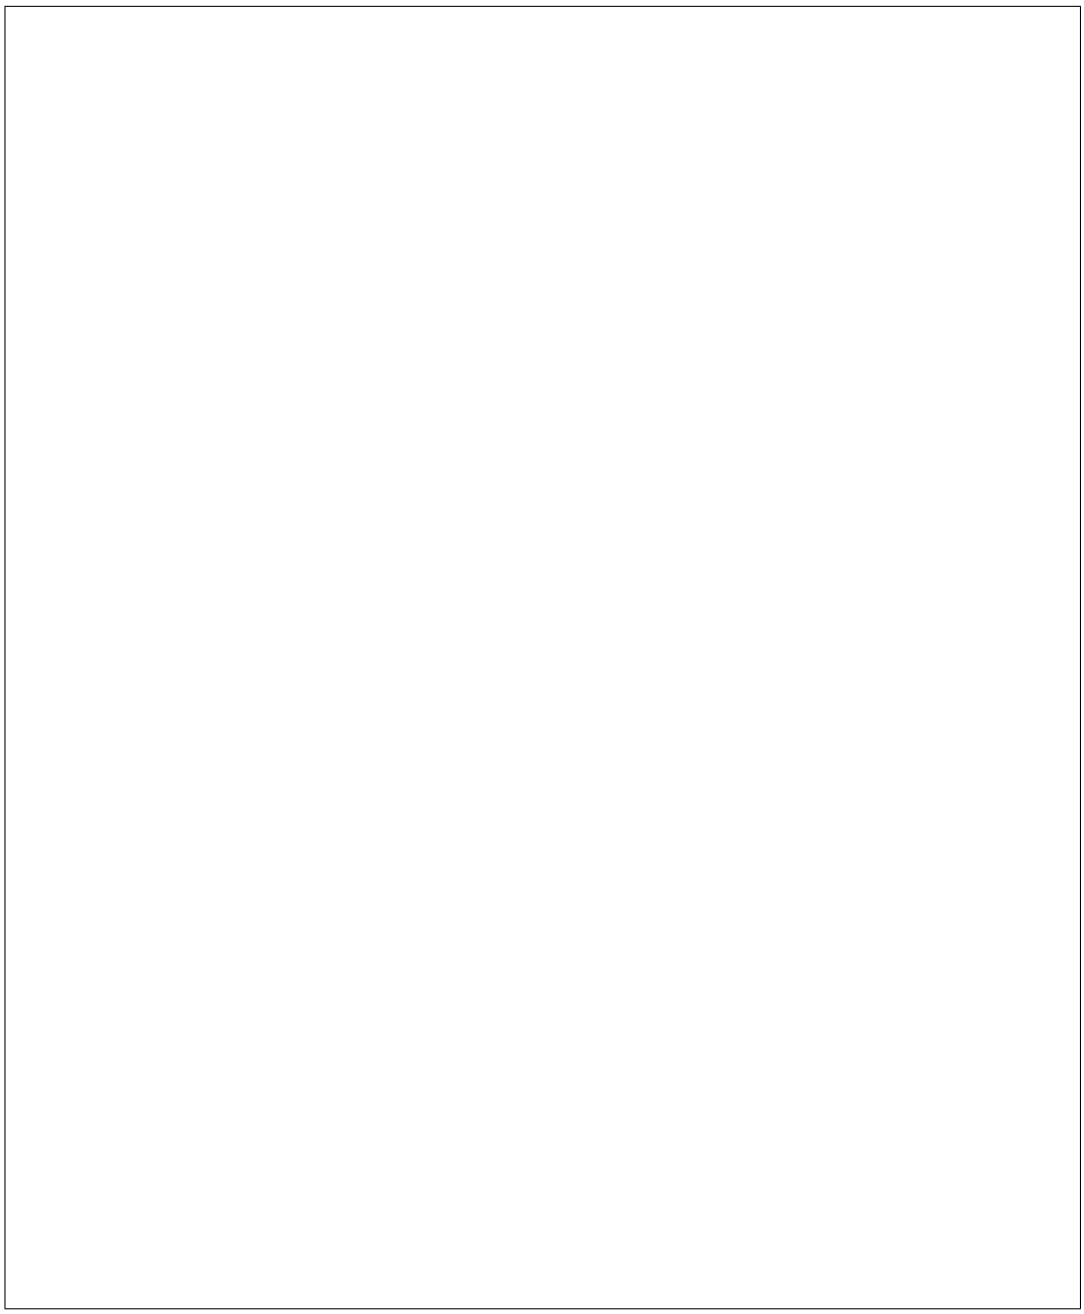

Indeed, the diagram on page 126 that outlines our logic can be extended directly to the three dimensional case. The only difference is the condition which will serve as our test for conservative fields, namely $\text{curl}\boldsymbol{F} = \mathbf{0}$.

The proofs of the links in the diagram for the three dimensional case below are very similar to those used in the two dimensional case. The only detail that is significantly different is showing that if $\text{curl}\boldsymbol{F} = \boldsymbol{0}$ then $\oint_C \boldsymbol{F} \cdot \boldsymbol{dr} = 0$. Note also that $\boldsymbol{F}$ must be a vector field defined everywhere in $\mathbb{R}^3$ with continuous partial derivatives. The proof of that part of the diagram requires a generalisation of Green's theorem known as *Stokes' theorem*, which we will investigate in the next section.

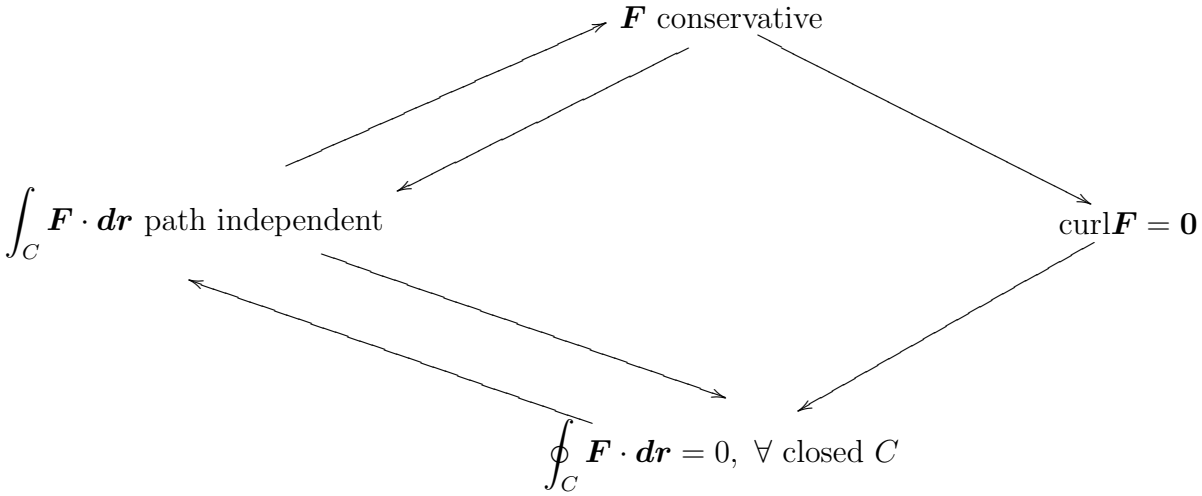

The main consequence of this diagram is that we have the following test for a conservative vector field in three dimensions:

A vector field $\boldsymbol{F}$ is conservative if and only if $\text{curl}\boldsymbol{F} = \boldsymbol{0}$.

26.3.1 Determine whether or not the vector field $\boldsymbol{F} = (1 + yz)\boldsymbol{i} + (1 + xz)\boldsymbol{j} + xy\boldsymbol{k}$ is conservative.

27 Stokes' theorem

By the end of this section, you should be able to answer the following questions:

- What is Stokes' theorem and under what conditions can it be applied?
- How do you apply Stokes' theorem?
- What is the circulation of a vector field?

27.1 Summary of surfaces and curves

Here we summarise the different types of curves and surfaces which we need to understand Stokes' theorem. Although most of these definitions have already been given, you may find it useful to have all of this information in one place so you can review at a glance.

27.1.1 Surfaces

- *Smooth*: the surface normal vector depends continuously on the points on the surface.
- *Piecewise smooth*: the surface consists of finitely many smooth surfaces intersecting only at their boundaries.
- *Oriented* (or *orientable*): the direction of the positive normal vector can be continued uniquely and continuously across the whole surface (especially if the surface is piecewise smooth).

27.1.2 Curves

- *Smooth*: the tangent at each point on the curve is unique and varies continuously.
- *Piecewise smooth*: the curve consists of finitely many smooth curves.
- *Simple*: the curve never intersects itself anywhere between its endpoints.

27.2 Stokes' theorem

Let S be a piecewise smooth, orientable surface in $\mathbb{R}^3$ and let the boundary of S be a piecewise smooth, simple, closed curve C. Let $\boldsymbol{F}(x, y, z)$ be a continuous vector function with continuous first partial derivatives in some domain containing S. Then

$$\iint_S (\text{curl}\boldsymbol{F}) \cdot \boldsymbol{n} \, dA = \oint_C \boldsymbol{F} \cdot \boldsymbol{dr},$$

where $\boldsymbol{n}$ is a unit normal vector of S, and the integration around C is taken in the direction using the "right hand rule" with $\boldsymbol{n}$.

27.2.1 Relation to Green's theorem

Recall Green's theorem in the plane. It relates a line integral on a boundary to a double integral over a region in the plane. Roughly speaking, Stokes' theorem is a 3-D version of this: it relates a surface integral on a piece of surface (in 3-D) to a line integral on the boundary of the surface.

In fact, note that if the surface is in the x-y plane with $\boldsymbol{n} = \boldsymbol{k}$, Stokes' theorem reduces to Green's theorem, since the $\boldsymbol{k}$ component of curl$\boldsymbol{F}$ is just $\dfrac{\partial F_2}{\partial x} - \dfrac{\partial F_1}{\partial y}$.

27.2.2 Verify Stokes' theorem where C is the curve of intersection of the plane $y + z = 2$ and the cylinder $x^2 + y^2 = 1$, oriented counterclockwise when looking from above, and $\boldsymbol{F} = [-y^2, x, z^2]$.

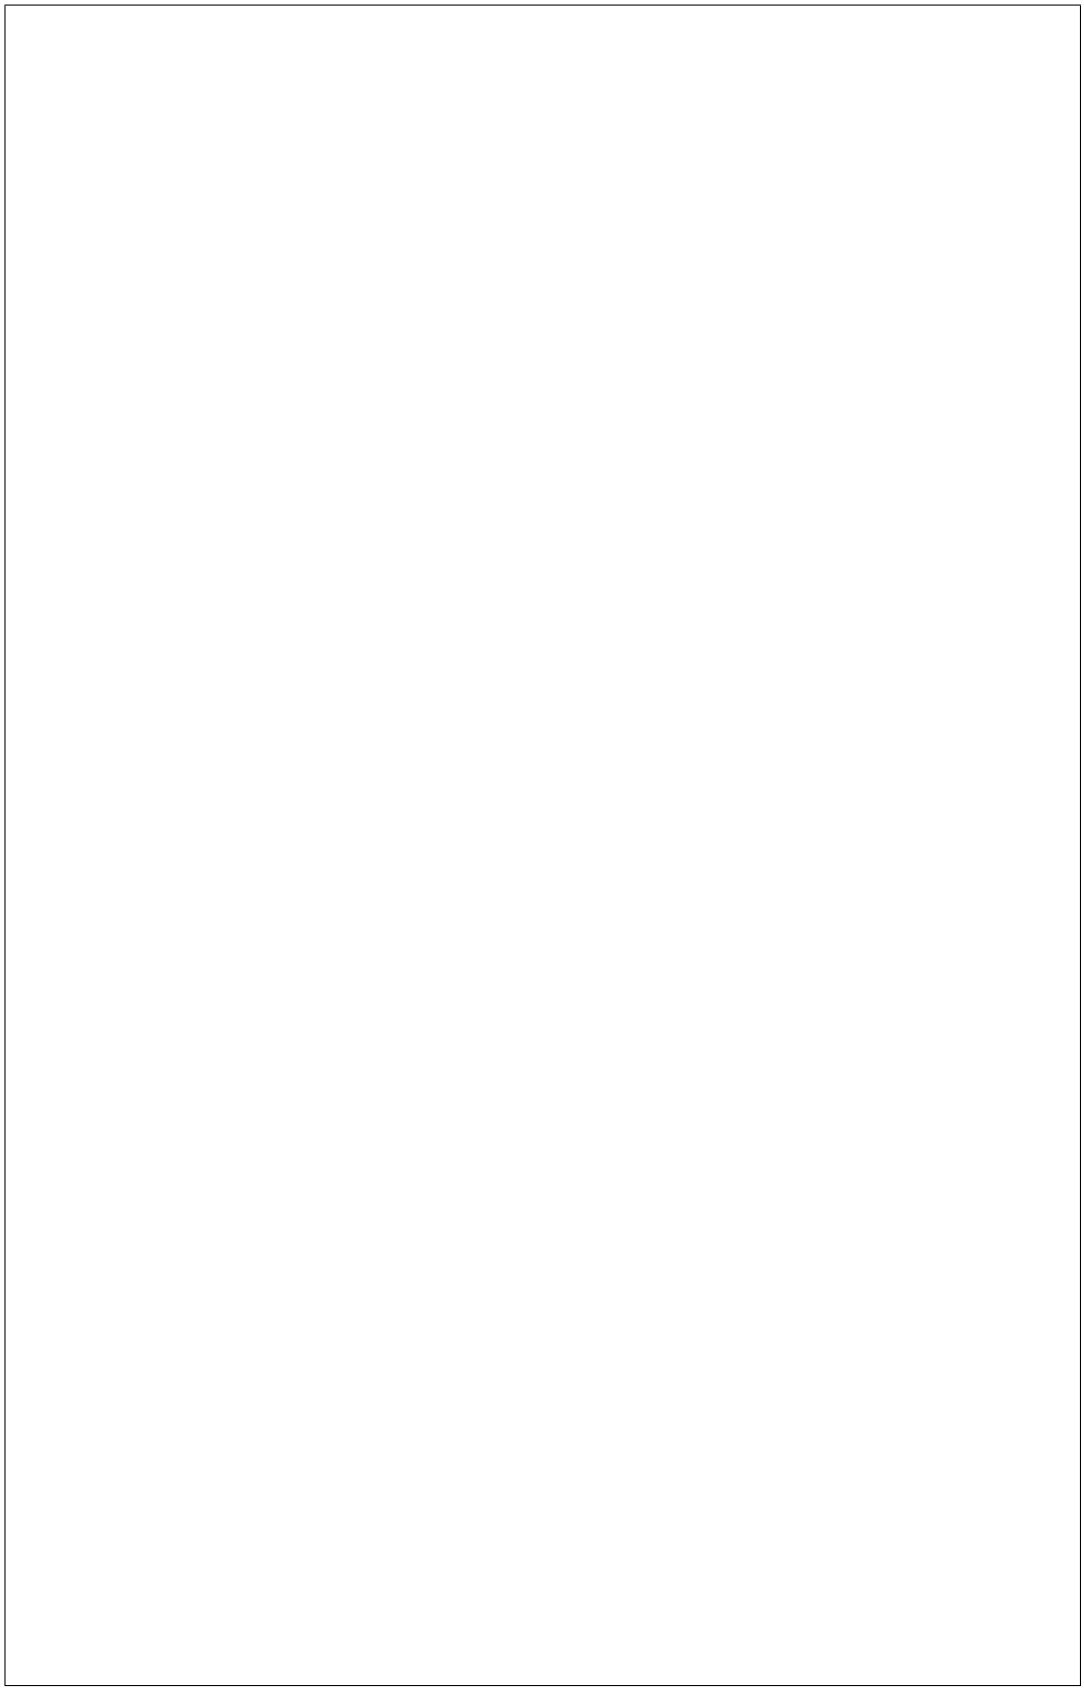

27.3 Circulation

Let $\boldsymbol{v}$ represent the velocity field of a fluid and C is a piecewise smooth, simple, closed curve. We have

$$\oint_C \boldsymbol{v} \cdot \boldsymbol{dr} = \oint_C \boldsymbol{v} \cdot \boldsymbol{T}\, dS,$$

where $\boldsymbol{T}$ is a unit tangent vector in the direction of the orientation of the curve. The dot product $\boldsymbol{v} \cdot \boldsymbol{T}$ is the component of $\boldsymbol{v}$ in the direction of $\boldsymbol{T}$ (and hence the curve), so we can interpret $\oint_C \boldsymbol{v} \cdot \boldsymbol{T}\, dS$ as a measure of the tendency of the fluid to move around the curve C. We call this quantity the *circulation* of $\boldsymbol{v}$ around C.

Now define a small circle C_a of radius a about a point P_0, such that the disc S_a enclosed by C_a is normal to the vector $\boldsymbol{n}(P_0)$. Our aim here is to better understand curl$\boldsymbol{v}$.

Since curl$\boldsymbol{v}$ is continuous, we approximate curl$\boldsymbol{v}$ over S_a as curl$\boldsymbol{v}(P_0)$. Stokes theorem then gives us

$$\begin{aligned}
\oint_{C_a} \boldsymbol{v} \cdot \boldsymbol{dr} &= \iint_{S_a} \text{curl}\boldsymbol{v} \cdot \boldsymbol{n}\, dS \\
&\approx \iint_{S_a} \text{curl}\boldsymbol{v}(P_0) \cdot \boldsymbol{n}(P_0)\, dS \\
&= \text{curl}\boldsymbol{v}(P_0) \cdot \boldsymbol{n}(P_0) \iint_{S_a} dS \\
&= \text{curl}\boldsymbol{v}(P_0) \cdot \boldsymbol{n}(P_0)(\pi a^2) \\
\Rightarrow \quad \text{curl}\boldsymbol{v}(P_0) \cdot \boldsymbol{n}(P_0) &\approx \frac{1}{\pi a^2} \oint_{C_a} \boldsymbol{v} \cdot \boldsymbol{dr} \\
&\approx \frac{\text{circulation around disc}}{\text{area of disc}}.
\end{aligned}$$

This approximation improves as $a \to 0$. Indeed

$$\text{curl}\boldsymbol{v}(P_0) \cdot \boldsymbol{n}(P_0) = \lim_{a \to 0} \frac{1}{\pi a^2} \oint_{C_a} \boldsymbol{v} \cdot \boldsymbol{dr}.$$

Note that this has a maximum value when curl$\boldsymbol{v}(P_0)$ and $\boldsymbol{n}(P_0)$ have the same direction.

In particular, if we take $\boldsymbol{n}(P_0)$ to be each of the coordinate unit vectors $\boldsymbol{i}$, $\boldsymbol{j}$, $\boldsymbol{k}$, we have the following: The $\boldsymbol{i}$, $\boldsymbol{j}$, $\boldsymbol{k}$ components of curl$\boldsymbol{v}(P_0)$ give the *circulation density* at P_0 in planes normal to each of the $\boldsymbol{i}$, $\boldsymbol{j}$, $\boldsymbol{k}$ respectively. The magnitude of curl$\boldsymbol{v}(P_0)$ gives the maximum circulation density about P_0 in a plane normal to curl$\boldsymbol{v}(P_0)$.

27.4 Curl fields and vector potentials

One immediate consequence is that if there are two different surfaces S_1 and S_2 satisfying the criteria of Stokes' theorem, both with the same boundary curve C, then

$$\iint_{S_1} \text{curl}\boldsymbol{F} \cdot \boldsymbol{n}_1 \, dS = \oint_C \boldsymbol{F} \cdot \boldsymbol{dr} = \iint_{S_2} \text{curl}\boldsymbol{F} \cdot \boldsymbol{n}_2 \, dS.$$

We have that if S is a closed surface satisfying all of the other criteria of Stokes' theorem, and if we define C to be any closed curve lying on S, so that S_1 and S_2 are two open surfaces whose union makes up S and whose common boundary is C, then

$$\begin{aligned} \oiint_S \text{curl}\boldsymbol{F} \cdot \boldsymbol{n} \, dS &= \iint_{S_1} \text{curl}\boldsymbol{F} \cdot \boldsymbol{n}_1 \, dS + \iint_{S_2} \text{curl}\boldsymbol{F} \cdot \boldsymbol{n}_2 \, dS \\ &= \oint_C \boldsymbol{F} \cdot \boldsymbol{dr} + \oint_{-C} \boldsymbol{F} \cdot \boldsymbol{dr} \\ &= \oint_C \boldsymbol{F} \cdot \boldsymbol{dr} - \oint_C \boldsymbol{F} \cdot \boldsymbol{dr} = 0, \end{aligned}$$

since the orientation of C as a boundary to S_1 will be in the opposite direction to that of S_2.

Let $\boldsymbol{F}$ be a vector field satisfying $\boldsymbol{F} = \text{curl}\boldsymbol{G}$ for some vector field $\boldsymbol{G}$. We call $\boldsymbol{F}$ a *curl field* and $\boldsymbol{G}$ a corresponding *vector potential*.

The above result says that the net outward flux of a curl field across any closed surface is zero.

We can verify that $\text{div}(\text{curl}\boldsymbol{G}) = 0$ for any vector field $\boldsymbol{G}$. Consequently we should not be too surprised by the above result, since Gauss' divergence theorem says that

$$\oiint_S (\text{curl}\boldsymbol{G}) \cdot \boldsymbol{n} \, dS = \iiint_V \text{div}(\text{curl}\boldsymbol{G}) \, dV = 0.$$

In fact, it turns out that we have the following test for curl fields:

Let $\boldsymbol{F}$ be a vector field whose components and their partial derivatives are continuous. If every closed surface in the domain of $\boldsymbol{F}$ only encloses points which are also in the domain of $\boldsymbol{F}$, and if $\text{div}\boldsymbol{F} = 0$, then there exists some $\boldsymbol{G}$ such that $\boldsymbol{F} = \text{curl}\boldsymbol{G}$. That is, $\boldsymbol{F}$ is a curl field.

28 Gaussian elimination and linear equations

By the end of this section, you should be able to answer the following questions:

- How do you use Gaussian elimination to find the row echelon form of a matrix?
- What are the three cases for solutions to systems of linear equations?
- How do you solve a system of linear equations?
- What are elementary matrices and how do they relate to elementary row operations?

Say we have m linear equations in n variables:

$$\begin{aligned} a_{11}x_1 + a_{12}x_2 + \cdots + a_{1n}x_n &= b_1 \\ a_{21}x_1 + a_{22}x_2 + \cdots + a_{2n}x_n &= b_2 \\ &\vdots \\ a_{m1}x_1 + a_{m2}x_2 + \cdots + a_{mn}x_n &= b_m. \end{aligned}$$

We can write these equations in matrix form: $A\boldsymbol{x} = \boldsymbol{b}$.

$A = [a_{ij}]$ is the $m \times n$ coefficient matrix.

$\boldsymbol{x} = \begin{pmatrix} x_1 \\ \vdots \\ x_n \end{pmatrix}$ is the column vector of unknowns, and $\boldsymbol{b} = \begin{pmatrix} b_1 \\ \vdots \\ b_m \end{pmatrix}$ is the column vector of the right hand side.

Note: a_{ij}, $b_j \in \mathbb{R}$ or $\mathbb{C}$.

28.1 Gaussian Elimination

To solve $A\boldsymbol{x} = \boldsymbol{b}$:
write *augmented matrix*: $[A|\boldsymbol{b}]$.

1. Find the left-most non-zero column, say column j.
2. Interchange top row with another row if necessary, so top element of column j is non-zero. (The **pivot**.)
3. Subtract multiples of row 1 from all other rows so all entries in column j below the top are then 0.

4. Cover top row; repeat **1** above on rest of rows.
Continue until all rows are covered, or until only $00\ldots0$ rows remain.

The result is a triangular system, easily solved by *back substitution*: solve the last equation first, then 2nd last equation and so on.

28.1.1 Example

Use Gaussian elimination to solve:

$$\begin{aligned} x_3 - x_4 &= 2 \\ -9x_1 - 2x_2 + 6x_3 - 12x_4 &= -7 \\ 3x_1 + x_2 - 2x_3 + 4x_4 &= 2 \\ 2x_3 &= 6 \end{aligned}$$

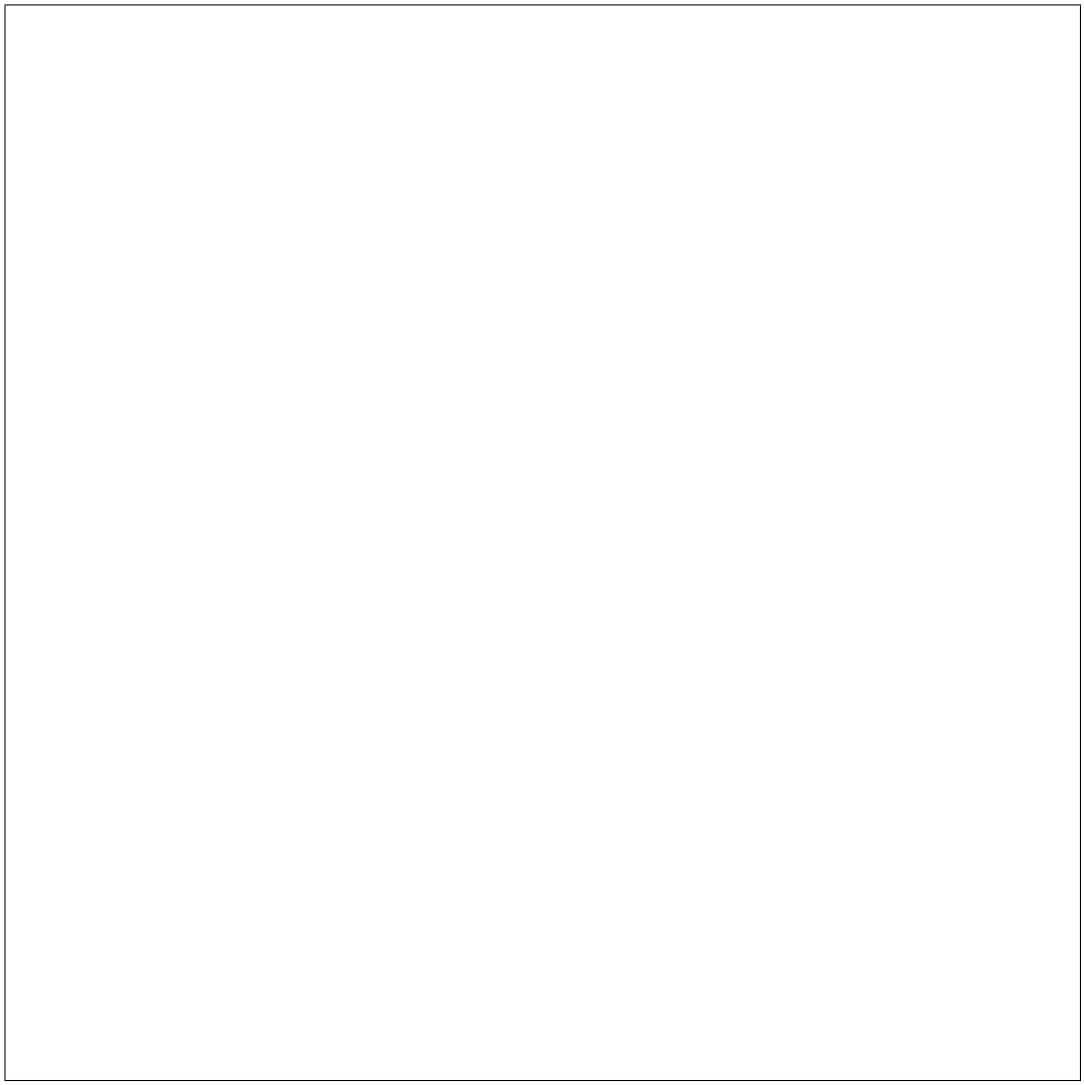

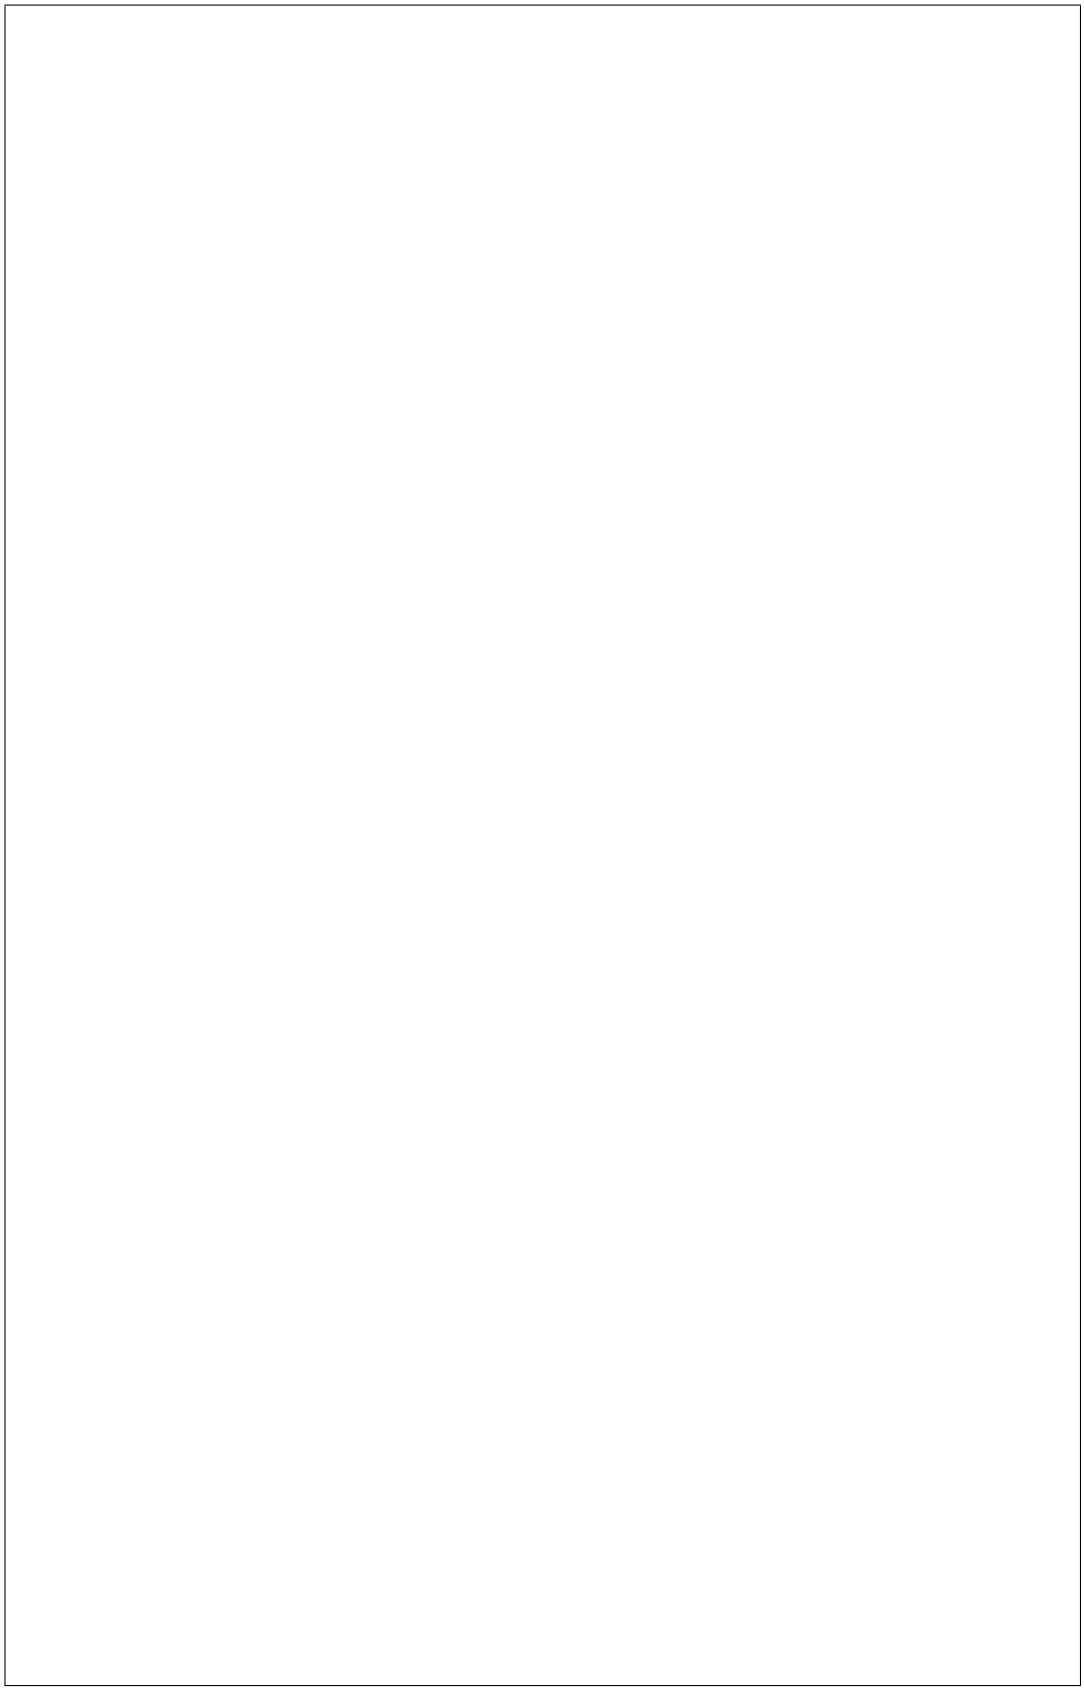

28.1.2 Definition (row echelon form)

A matrix is in *row echelon form* (r.e.f.) if each row after the first starts with *more* zeros than the previous row (or else rows at bottom of matrix are all zeros).

The Gauss algorithm converts any matrix to one in row echelon form. The 2 matrices are *equivalent*, that is, they have the same solution set.

28.1.3 Elementary row operations

1. $r_i \leftrightarrow r_j$: swap rows i and j.
2. $r_i \to r_i - cr_j$: replace row i with
(row i minus c times row j).
3. $r_i \to cr_i$:
replace row i with c times row i, where $c \neq 0$.

The Gauss algorithm uses only **1** and **2**.

28.2 Possible solutions for $A\boldsymbol{x} = \boldsymbol{b}$

Consider the r.e.f. of $[A|\boldsymbol{b}]$. Then we have three possibilities:

(1) *Exactly one* solution; here the r.e.f. gives each variable a single value, so the number of variables, n, equals the number of non-zero rows in the r.e.f.

(2) *No* solution; when one row of r.e.f. is $(0\,0\,\ldots\,d)$ with $d \neq 0$. We can't solve $0x_1 + 0x_2 + \cdots + 0x_m = d$ if $d \neq 0$; it says $0 = d$. In this case the system is said to be *inconsistent*.

(3) *Infinitely many* solutions; here the number of non-zero rows of the r.e.f. is *less* than the number of variables.

Note that a *homogeneous* system has $\boldsymbol{b} = \boldsymbol{0}$, i.e., all zero RHS. Then we always have at least the trivial solution, $x_i = 0$, $1 \leq i \leq n$.

28.2.1 Examples

$$\begin{aligned} x_1 + x_2 - \quad x_3 &= 0 \\ 2x_1 - x_2 &= 0 \\ 4x_1 + x_2 - 2x_3 &= 1 \end{aligned}$$

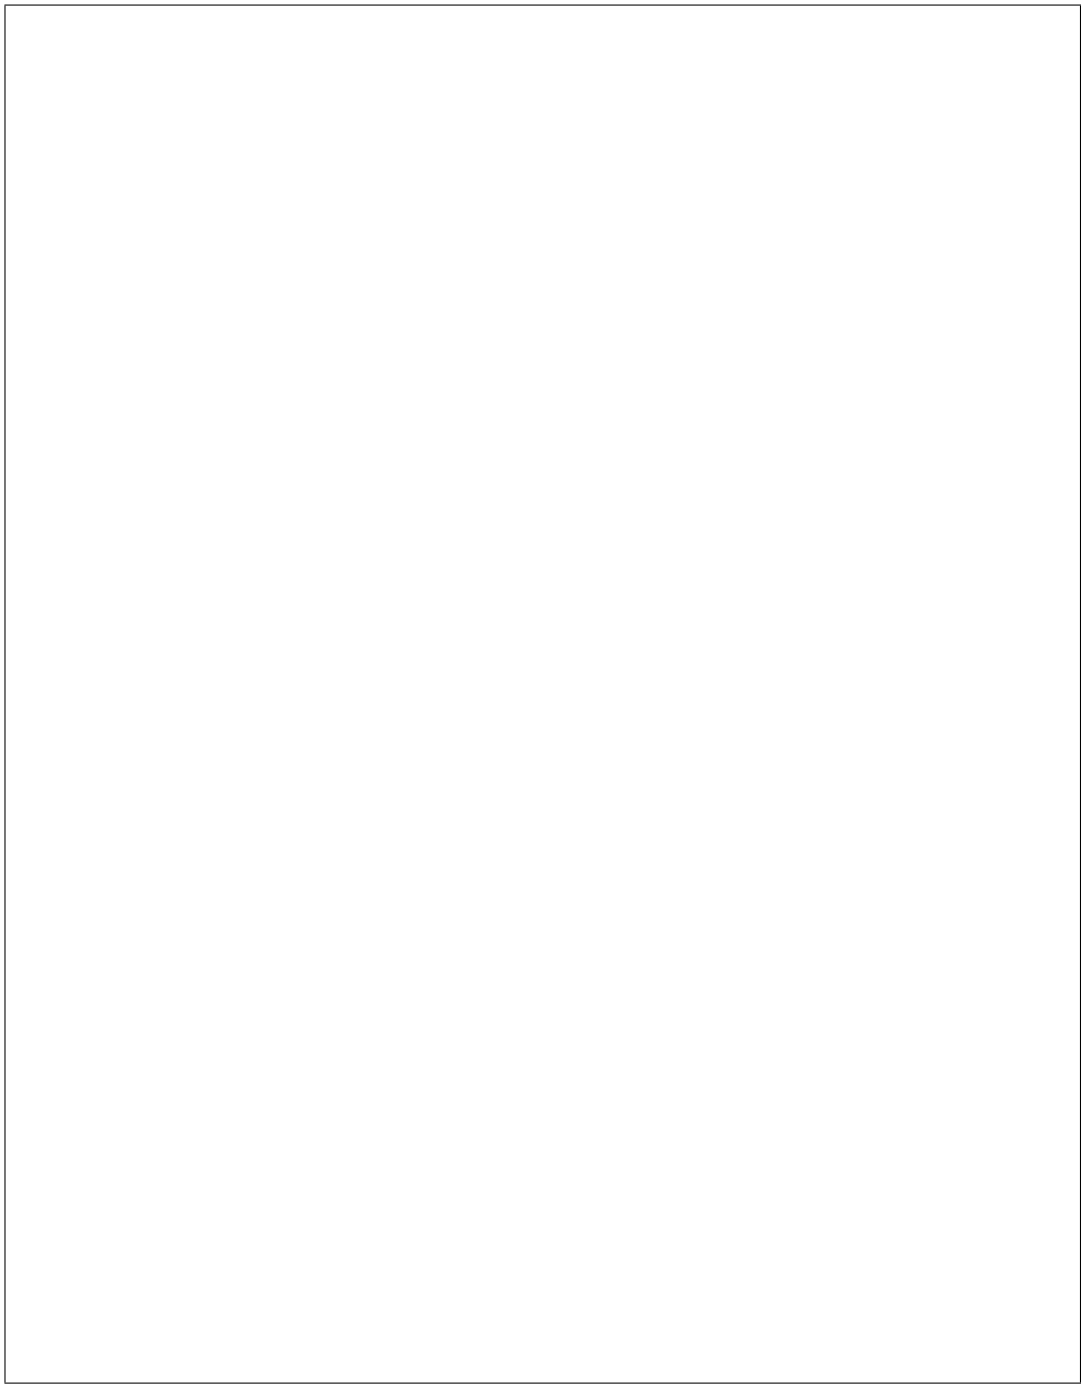

$$\begin{aligned} x_1 - 2x_2 + 4x_3 &= 2 \\ 2x_1 - 3x_2 + 7x_3 &= 6 \\ x_2 - x_3 &= 2 \end{aligned}$$

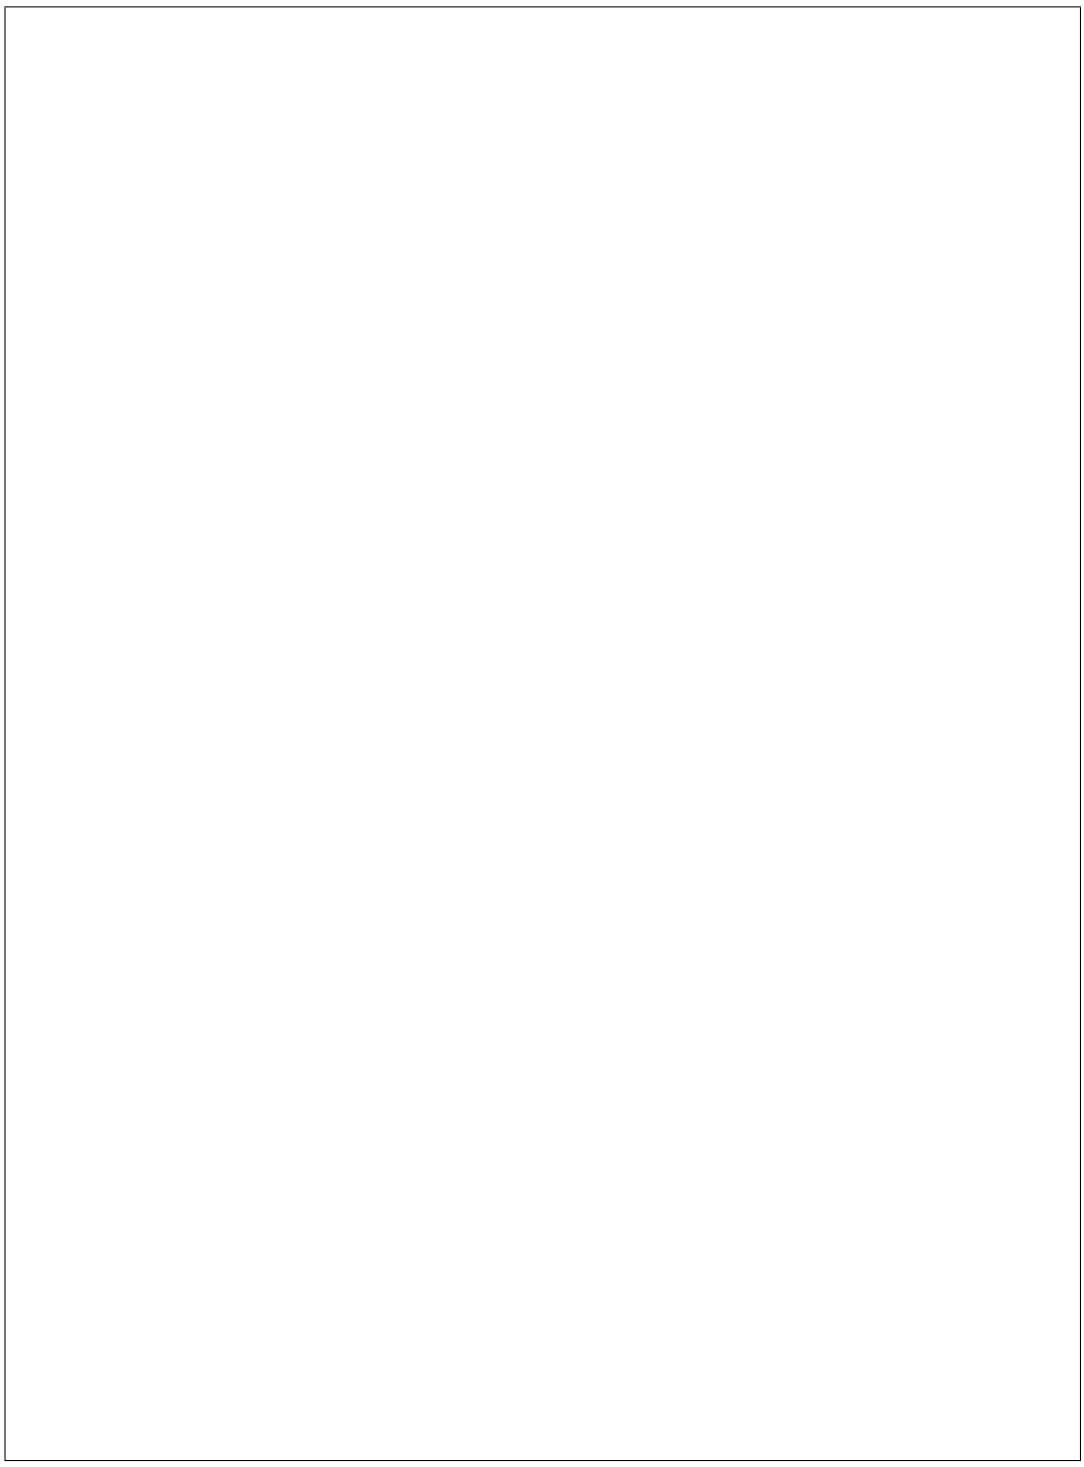

28.3 Elementary matrices

An $n \times n$ matrix is called elementary if it can be obtained from the $n \times n$ identity matrix by performing one of the three elementary row operations.

For example, for 3×3 matrices,

- A type 1 row operation is $R_2 \leftrightarrow R_3$ which corresponds to the elementary matrix
$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}.$$
- A type 2 row operation is $R_2 \to R_2 - 3R_1$ which corresponds to the elementary matrix
$$\begin{pmatrix} 1 & 0 & 0 \\ -3 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$
- A type 3 row operation is $R_3 \to 5R_3$ which corresponds to the elementary matrix
$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 5 \end{pmatrix}.$$

In fact, applying an elementary row operation to any $n \times m$ matrix A is equivalent to multiplying A from the left by the corresponding elementary matrix.

For example,

$$\begin{pmatrix} 1 & 0 & 2 & 3 \\ 2 & -1 & 3 & 6 \\ 1 & 4 & 4 & 0 \end{pmatrix} \overset{R_2 \leftrightarrow R_3}{\longrightarrow} \begin{pmatrix} 1 & 0 & 2 & 3 \\ 1 & 4 & 4 & 0 \\ 2 & -1 & 3 & 6 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 & 2 & 3 \\ 2 & -1 & 3 & 6 \\ 1 & 4 & 4 & 0 \end{pmatrix}$$

and

$$\begin{pmatrix} 1 & 0 & 2 & 3 \\ 2 & -1 & 3 & 6 \\ 1 & 4 & 4 & 0 \end{pmatrix} \overset{R_2 \to R_2 - 3R_1}{\longrightarrow} \begin{pmatrix} 1 & 0 & 2 & 3 \\ -1 & -1 & -3 & -3 \\ 1 & 4 & 4 & 0 \end{pmatrix}$$
$$= \begin{pmatrix} 1 & 0 & 0 \\ -3 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 2 & 3 \\ 2 & -1 & 3 & 6 \\ 1 & 4 & 4 & 0 \end{pmatrix}.$$

Elementary matrices are useful theoretical tools. Many proofs of fundamental results in linear algebra rely on these matrices and their properties.

For example, we can view the steps in determining the inverse of a square matrix as a sequence of operations involving elementary matrices. We look for a solution X to the matrix equation $AX = I$ by forming the augmented matrix $(A|I)$ and performing elementary row operations. For example, after performing three operations, we have really changed the equation to

$$E_3E_2E_1AX = E_3E_2E_1I,$$

where E_1, E_2, E_3 are elementary matrices. On completion of the steps (say there are n of them), we reach

$$E_n \dots E_2E_1AX = IX = X = E_n \dots E_2E_1I,$$

which tells us that the inverse of A (usually denoted A^{-1}), if it exists, is nothing more than the product of elementary matrices $E_n \dots E_2E_1$.

In fact, if $A^{-1} = E_n \dots E_2E_1$, then the matrix A itself must be a product of inverses of elementary matrices $A = E_1^{-1}E_2^{-1} \dots E_n^{-1}$. It turns out, as we shall see, that the inverse of an elementary matrix is an elementary matrix. Hence if A is invertible, then it can be written as a product of elementary matrices.

28.3.1 Two important results regarding determinants

Two significant results regarding determinants are

$$\det(AB) = \det(A)\det(B), \quad \det(A) = \det(A^T)$$

You should already be familiar with these results. They can be proved by the use of elementary matrices, by first establishing the results where A is an elementary matrix, and then generalising.

The proofs are beyond the scope of this course, but it is worth mentioning that the proofs make use of elementary matrices, hence demonstrating their importance.

28.3.2 Inverses of elementary matrices

It is a simple matter to verify that elementary matrices of type 1 (corresponding to the row operation of swapping two rows) square to the identity. In other words, the inverses of these matrices are just the matrices themselves. See, for example, that

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Note that in section 30.1 we refer to these elementary matrices as permutation matrices.

Let us consider elementary matrices of type 2. It is straightforward to give the inverse of these matrices. Note the pattern for the following 3×3 matrices:

$$\begin{pmatrix} 1 & 0 & 0 \\ a & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ -a & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ b & 0 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -b & 0 & 1 \end{pmatrix},$$

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

Finally, the inverses of the type 3 elementary matrices are simply

$$\begin{pmatrix} a & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} 1/a & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \begin{pmatrix} 1 & 0 & 0 \\ 0 & b & 0 \\ 0 & 0 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1/b & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & c \end{pmatrix}^{-1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1/c \end{pmatrix}.$$

The significance here is that in general, the inverse of an elementary matrix is an elementary matrix of the same type.

Say we need to perform three elementary row operations to obtain a r.e.f. of A. We can then write

$$E_3E_2E_1A = U,$$

where U is the r.e.f. of A. Since we know the inverses of all elementary matrices (indeed, they do exist), we can write

$$A = E_1^{-1}E_2^{-1}E_3^{-1}U.$$

By observation, the matrix $L = E_1^{-1}E_2^{-1}E_3^{-1}$ is lower triangular with 1's on the main diagonal.

For the matrix $A = \begin{pmatrix} 1 & -1 & 4 \\ 1 & 0 & -2 \\ 2 & -2 & 10 \end{pmatrix}$, the two operations which give the r.e.f. are $R_2 \to R_2 - R_1$ and $R_3 \to R_3 - 2R_1$, so the r.e.f of A can be expressed in terms of elementary matrices

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -2 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} A.$$

Since we can easily invert these elementary matrices, we have

$$A = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix} U$$

$$\Rightarrow \quad A = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 2 & 0 & 1 \end{pmatrix} U.$$

29 LU decompositions

By the end of this section, you should be able to answer the following questions:

- How do you find an LU decomposition of a matrix?
- How do you use an LU decomposition to solve a system of equations?

29.1 Finding L and U

Given an $m \times n$ matrix A, we use the Gauss algorithm to find the r.e.f. U (which is also $m \times n$) for A.

Say *no* row interchanges are used, so there are only operations of the form $r_i \to r_i - cr_j$. Let c_{i1} be the multiple of the 1st row subtracted from the ith row, c_{i2} be the multiple of the 2nd row subtracted from the (new) ith row, etc., when finding U.
Form the $m \times m$ lower triangular matrix:

$$L = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 \\ c_{21} & 1 & 0 & \dots & 0 \\ c_{31} & c_{32} & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ c_{m1} & c_{m2} & c_{m3} & \dots & 1 \end{pmatrix}$$

Our main result is that indeed $A = LU$. This is what we call an LU decomposition of A.

It is worth stressing that *it is only possible to find an LU decomposition if no row interchanges are used.*

We remark that not every matrix has an LU decomposition. If, however, a matrix does have an LU decomposition, then

$$\det A = \det L \det U = \det U.$$

29.1.1 Example

Find the LU decomposition for A where $A = \begin{pmatrix} 2 & 3 & -1 \\ 4 & 5 & 0 \\ -2 & -6 & 8 \end{pmatrix}$, then calculate $\det A$.

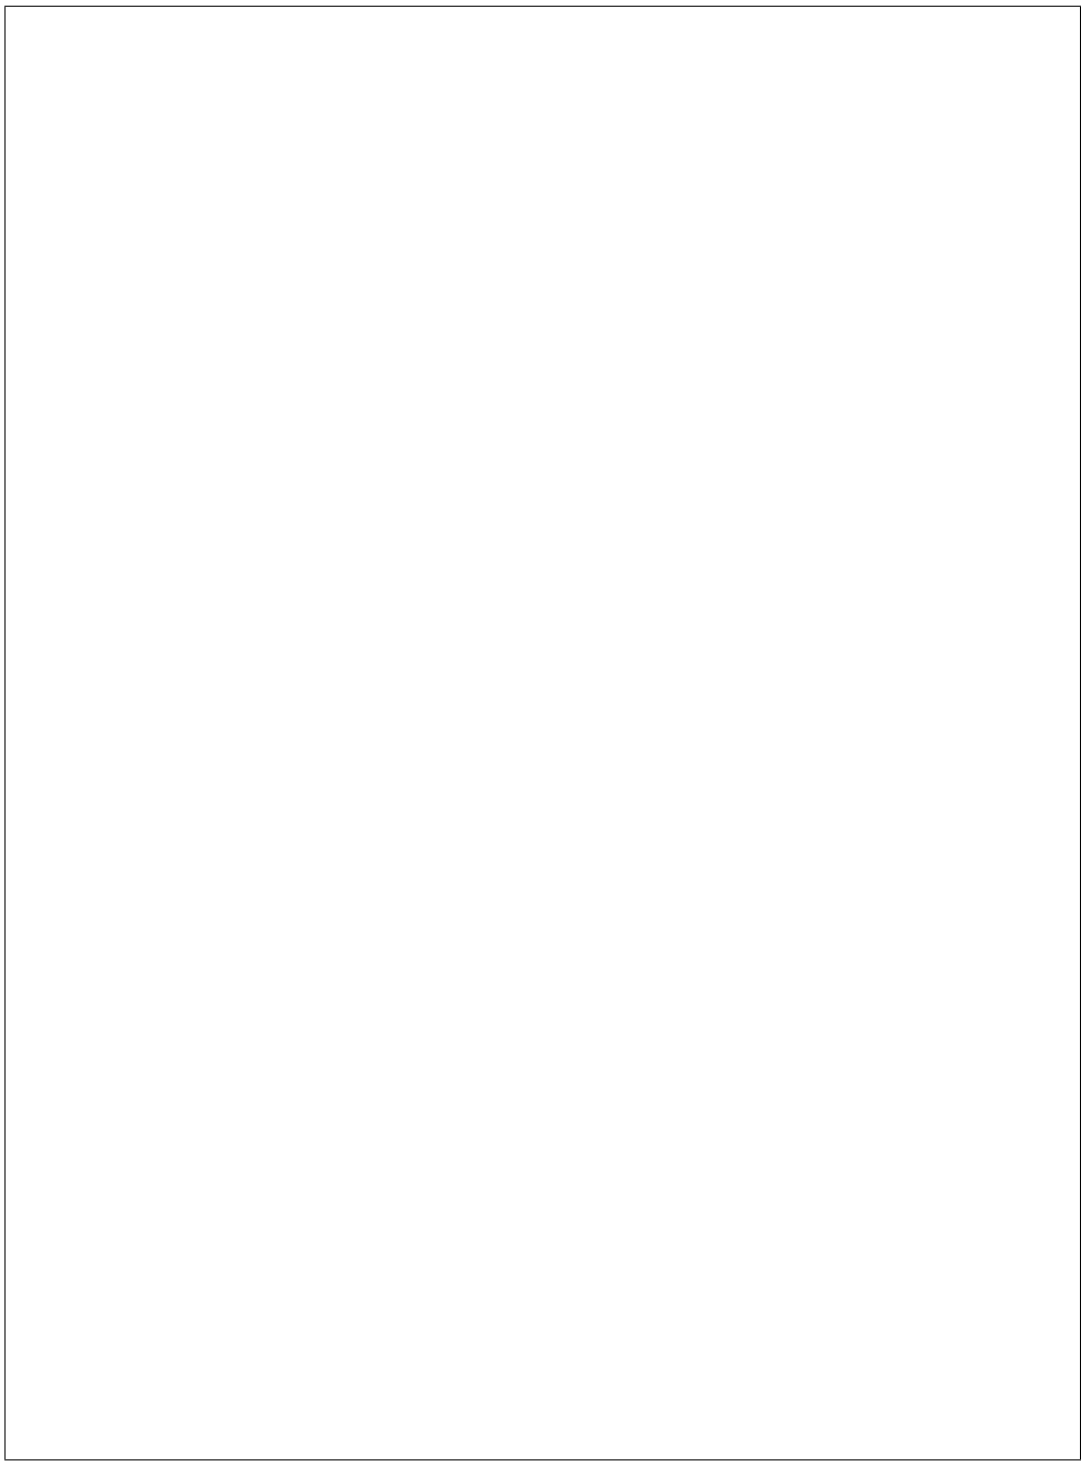

29.2 Using an LU decomposition to solve systems of equations

We can use this decomposition to solve $A\boldsymbol{x} = \boldsymbol{b}$ by first setting $\boldsymbol{y} = U\boldsymbol{x}$ and then solving $L\boldsymbol{y} = \boldsymbol{b}$ to obtain $\boldsymbol{y}$, and then solving $U\boldsymbol{x} = \boldsymbol{y}$ to obtain the solution $\boldsymbol{x}$.

Since L is lower triangular and U is in r.e.f., solving $L\boldsymbol{y} = \boldsymbol{b}$ (by forward substitution) and $U\boldsymbol{x} = \boldsymbol{y}$ (by back substitution) are both straightforward.

The advantage of this method is that we only need to compute L and U once. Then we can use them for many different $\boldsymbol{b}$, even when perhaps $\boldsymbol{b}_j$ depends upon earlier $\boldsymbol{b}$. This method also works if A is singular.

29.2.1 Example

Given $A = \begin{pmatrix} 2 & 3 & -1 \\ 4 & 5 & 0 \\ -2 & -6 & 8 \end{pmatrix}$, solve $A\boldsymbol{x} = \begin{pmatrix} 1 \\ 2 \\ 3 \end{pmatrix}$.

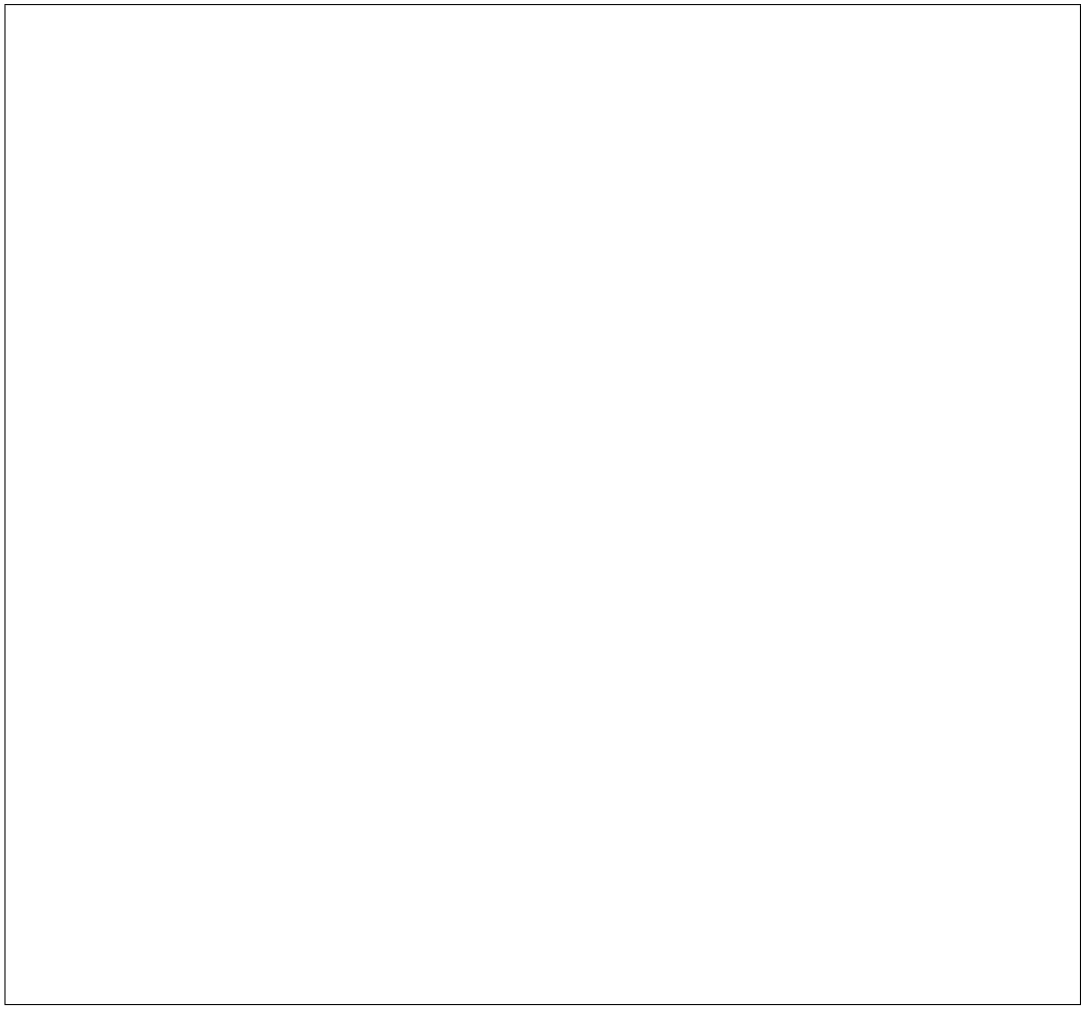

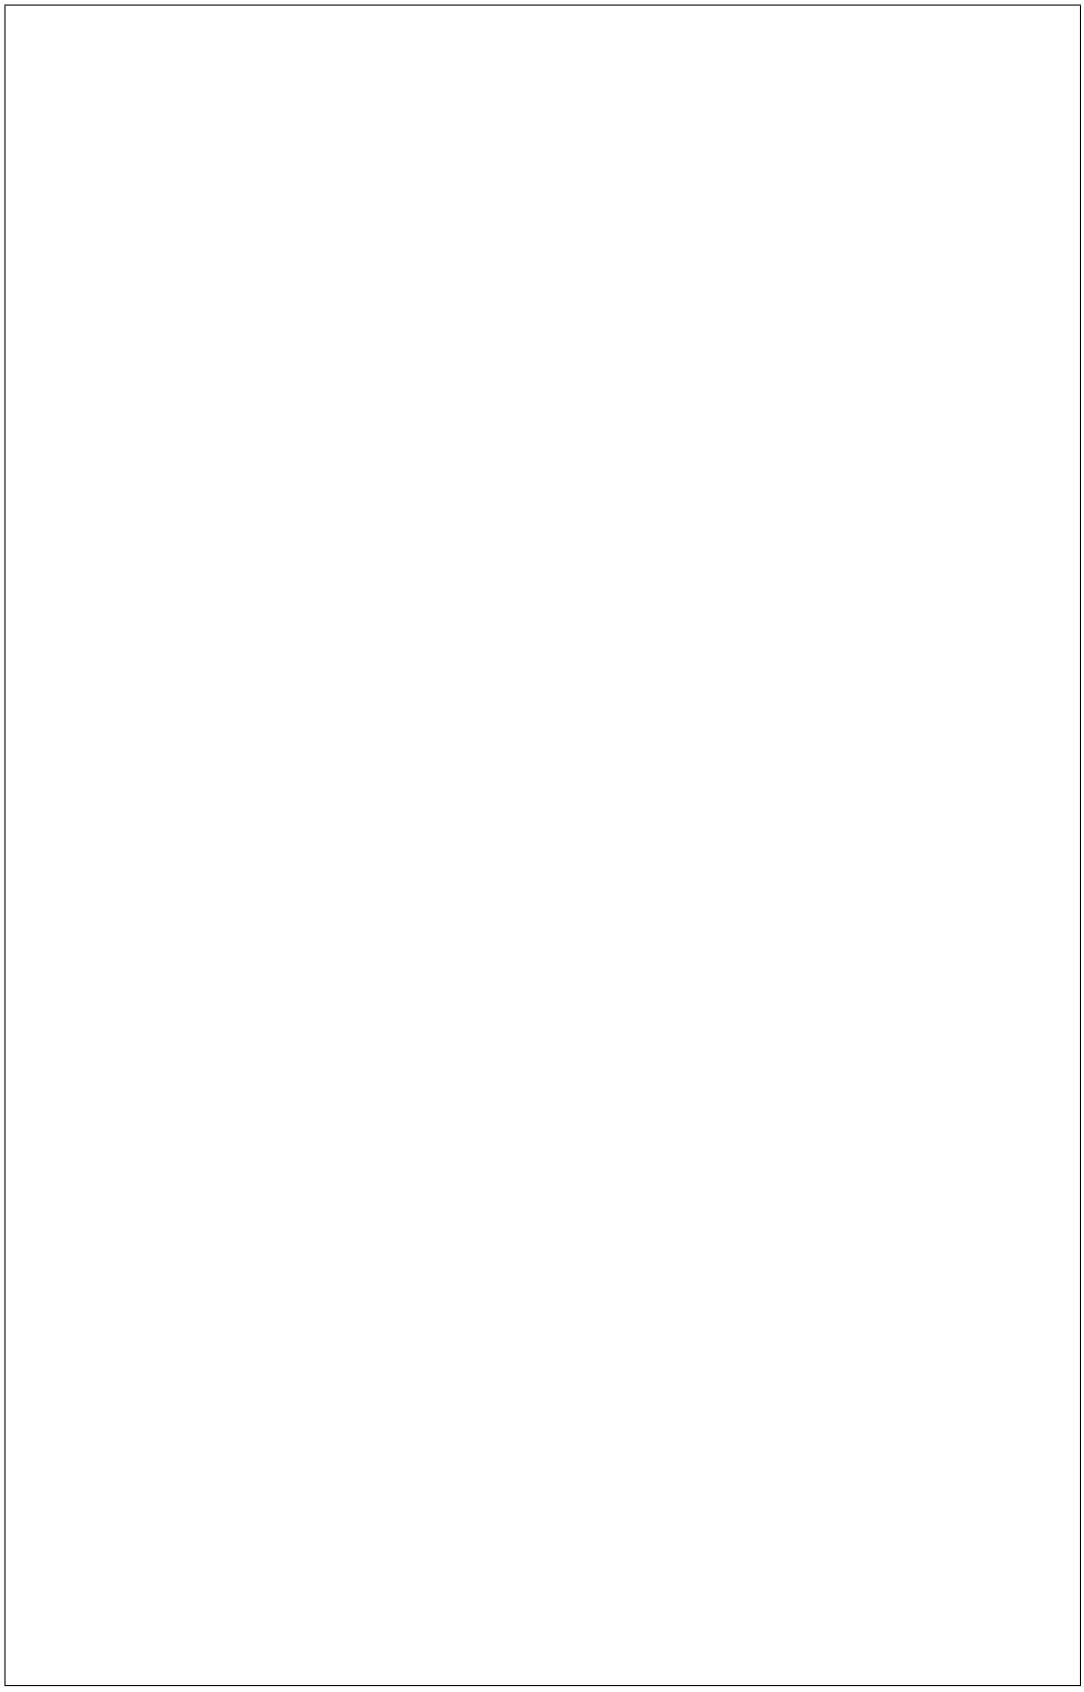

30 Permutation matrices and PLU decompositions

By the end of this section, you should be able to answer the following questions:

- How do you find a PLU decomposition of a matrix?
- How do you use a PLU decomposition to calculate a matrix determinant?

We mentioned in the last section that we can only find an LU decomposition if no row interchanges are needed to obtain the r.e.f. of a matrix. What if we *do* need row interchanges to get the r.e.f.?

30.1 Definition of permutation matrix

A *permutation matrix* is a matrix obtained from an identity matrix I by interchanging any 2 rows.

Define $P_{k,\ell}^{(n)}$ as the permutation matrix obtained from the $n \times n$ identity I by swapping rows k and ℓ.

So, for example

$$P_{2,3}^{(3)} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$$

and

$$P_{2,3}^{(4)} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Note that $P_{i,j}^{(k)} P_{i,j}^{(k)} = I^{(k)}$.

If A is $m \times n$, then $P_{k,\ell}^{(m)} A$ is a matrix obtained from A by swapping its rows k and ℓ.

If row interchanges are needed to get r.e.f. U from A, we could *first* rearrange all rows of A so that *no* interchanges are subsequently needed.
Say the system $A\boldsymbol{x} = \boldsymbol{b}$ is replaced by $A'\boldsymbol{x} = \boldsymbol{b}'$ after a series of row swaps, such that $A' = LU$. Then $A = PA'$ where P is a product of permutation matrices (maybe several).
Hence $A = PLU$.

30.2 Theorem (PLU decomposition)

Every $m \times n$ matrix A can be written in the form $A = PLU$ where P is a product of permutation matrices, L is an $m \times m$ lower triangular matrix with its main diagonal entries all 1, and U is an $m \times n$ r.e.f. matrix.

30.2.1 Example

Find a PLU decomposition of $A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 2 & 4 \end{pmatrix}$.

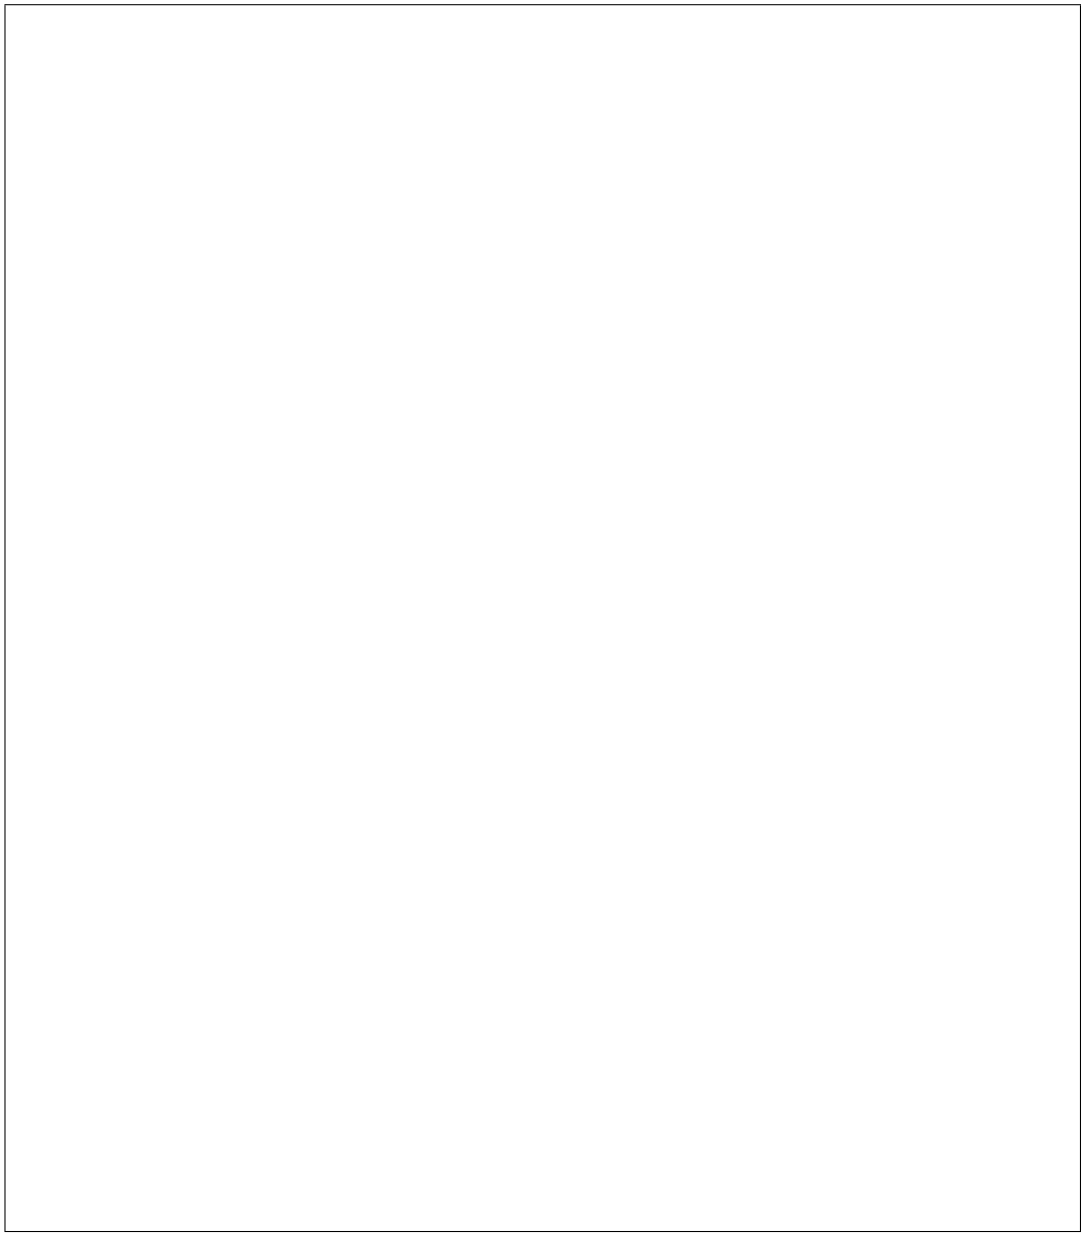

30.3 Determinants

This also gives an efficient way to find $\det(A)$, for a *square* matrix A.
If U is r.e.f. for A, found by using the Gauss algorithm on A, then

$$\det(A) = (-1)^N \det(U),$$

where N is the number of row interchanges used.

For: $\det(XY) = \det(X)\det(Y)$, so $\det(A) = \det(PLU) = \det(PL)\det(U)$.
But L is lower triangular with all 1s on main diagonal so $\det(L) = 1$.
PL is L with various rows interchanged.
Interchanging two rows of any determinant changes its sign. Hence $\det(PL) = \pm\det(L) = \pm 1$.

30.3.1 Example

Find a PLU decomposition of the matrix $A = \begin{pmatrix} 2 & 3 & -1 & 2 \\ -4 & -6 & 2 & 1 \\ 2 & 4 & 4 & -1 \\ 4 & 8 & 2 & 7 \end{pmatrix}$, then calculate its determinant. Note that it is not obvious in this case which rows to swap, so we treat it like a normal LU decomposition and then swap rows if required.

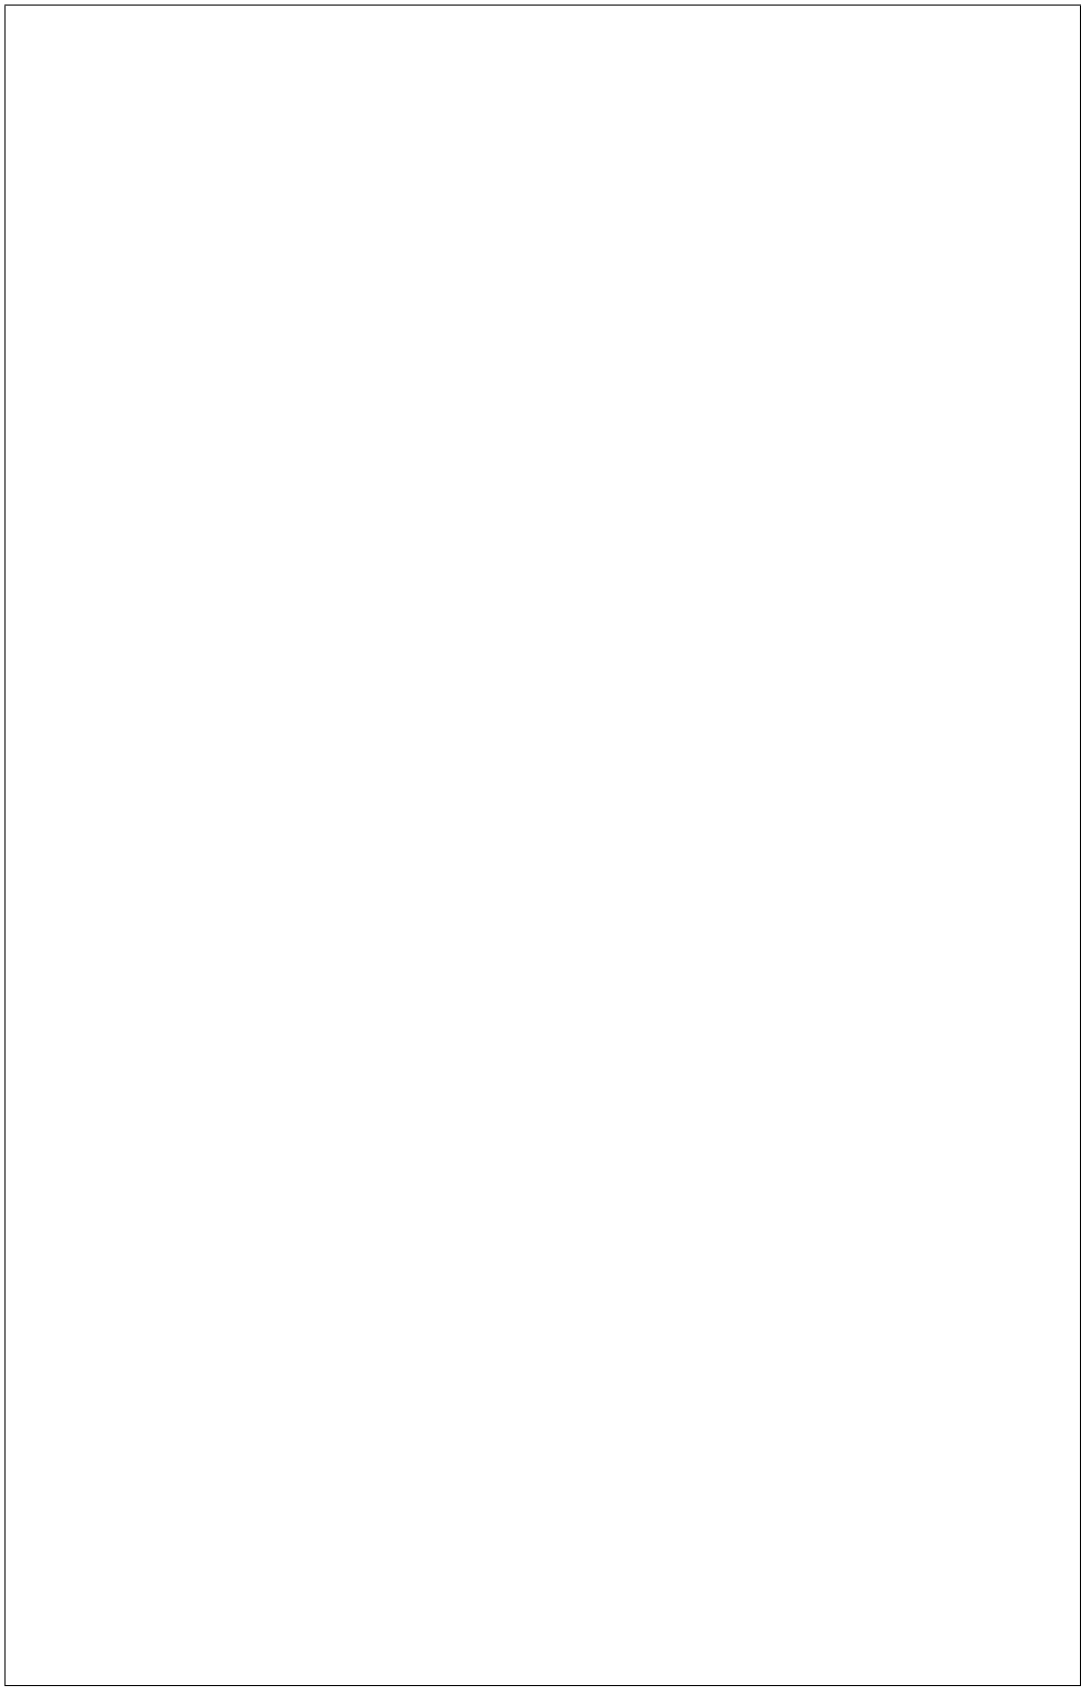

31 Eigenvalues and eigenvectors

By the end of this section, you should be able to answer the following questions:

- How do you find the eigenvalues and eigenvectors of a given square matrix?
- What are some simple properties of eigenvalues and eigenvectors?
- Prove that the eigenvectors corresponding to distinct eigenvalues are linearly independent.

A great deal of this section should be familiar to you. We start by recalling some results on vector spaces associated with matrices.

31.1 Column space, row space, rank, nullity

For any $m \times n$ real matrix A, the *null space* of A is the vector space

$$N(A) = \{\boldsymbol{x} \in \mathbb{R}^n \mid A\boldsymbol{x} = \boldsymbol{0}\}.$$

The dimension of $N(A)$ is called the *nullity* of A, denoted nullity(A).

The column space of an $m \times n$ matrix A is the space spanned by the column vectors of A ($\subseteq \mathbb{R}^m$). The dimension of the column space of A is called the *rank* of A, denoted rank(A). This coincides with the number of non-zero rows in the r.e.f. of A.

The row space of an $m \times n$ matrix A is the space spanned by the row vectors of A ($\subseteq \mathbb{R}^n$). A basis is given by the non-zero rows in the r.e.f. of A. The dimension of the row space is also given by the rank of A.

Note that the row space of A^T = column space of A.

For $m \times n$ matrices,

$$\text{rank}(A) + \text{nullity}(A) = n.$$

31.2 Non-singular matrices

For $n \times n$ square matrix A, we have several conditions for the existence of A^{-1}.

For $n \times n$ matrix A, the following are equivalent:
1. A is non-singular.
2. $A\boldsymbol{x} = \mathbf{0}$ has only the trivial solution $\boldsymbol{x} = \mathbf{0}$.
3. If U is a r.e.f. for A, then U has no row of all zeros.
4. $A\boldsymbol{x} = \boldsymbol{b}$ has a solution for every n-dimensional column vector $\boldsymbol{b}$.
5. $\det(A) \neq 0$.
6. The columns of A are linearly independent.
7. The rows of A are linearly independent.
8. $\text{nullity}(A) = 0$.
9. $\text{rank}(A) = \text{n}$.

31.3 Eigenvalues and eigenvectors

Let A be a square matrix. Then an *eigenvector* of A is a vector $\boldsymbol{v} \neq \mathbf{0}$ such that

$$A\boldsymbol{v} = \lambda\boldsymbol{v},$$

for some scalar λ.
The scalar λ is called the corresponding *eigenvalue*.

If $\boldsymbol{v}$ is an eigenvector of A, then so is $t\boldsymbol{v}$ for any scalar $t \neq 0$.

Recall if λ is an eigenvalue of A, with corresponding eigenvector $\boldsymbol{v}$, then $A\boldsymbol{v} = \lambda\boldsymbol{v} = \lambda I\boldsymbol{v}$, so $(A - \lambda I)\boldsymbol{v} = \mathbf{0}$. Hence $\boldsymbol{x} = \boldsymbol{v}$ is a non-trivial solution to the homogeneous system of equations $(A - \lambda I)\boldsymbol{x} = \mathbf{0}$, and conversely, if there is a non-trivial solution then λ is an eigenvalue of A. Thus:

λ is an eigenvalue of A
if and only if $(A - \lambda I)\boldsymbol{x} = \mathbf{0}$ has a non-trivial solution
if and only if $A - \lambda I$ is singular
if and only if $\det(A - \lambda I) = 0$.

For an $n \times n$ matrix A, $\det(A - \lambda I)$ is a polynomial of degree n in λ, called the *characteristic polynomial* of A.

The equation $\det(A - \lambda I) = 0$ is the *characteristic equation* of A.

Eigenvalues λ may be complex numbers, and the eigenvectors $\boldsymbol{v}$ may have complex components, even for real matrices A.

To find the eigenvalues and eigenvectors, do the following:

1. Find the roots of the characteristic polynomial, $\det(A - \lambda I) = 0$. These are the eigenvalues.
2. For each eigenvalue λ, find all $\boldsymbol{v}$ satisfying $(A - \lambda I)\boldsymbol{v} = \mathbf{0}$. These are the eigenvectors. The vector space spanned by the eigenvectors corresponding to each eigenvalue is called the *eigenspace* associated to λ.

31.3.1 Example

Find the eigenvalues and eigenvectors of $A = \begin{pmatrix} -3 & 1 & 0 \\ 1 & -2 & 1 \\ 0 & 1 & -3 \end{pmatrix}$.

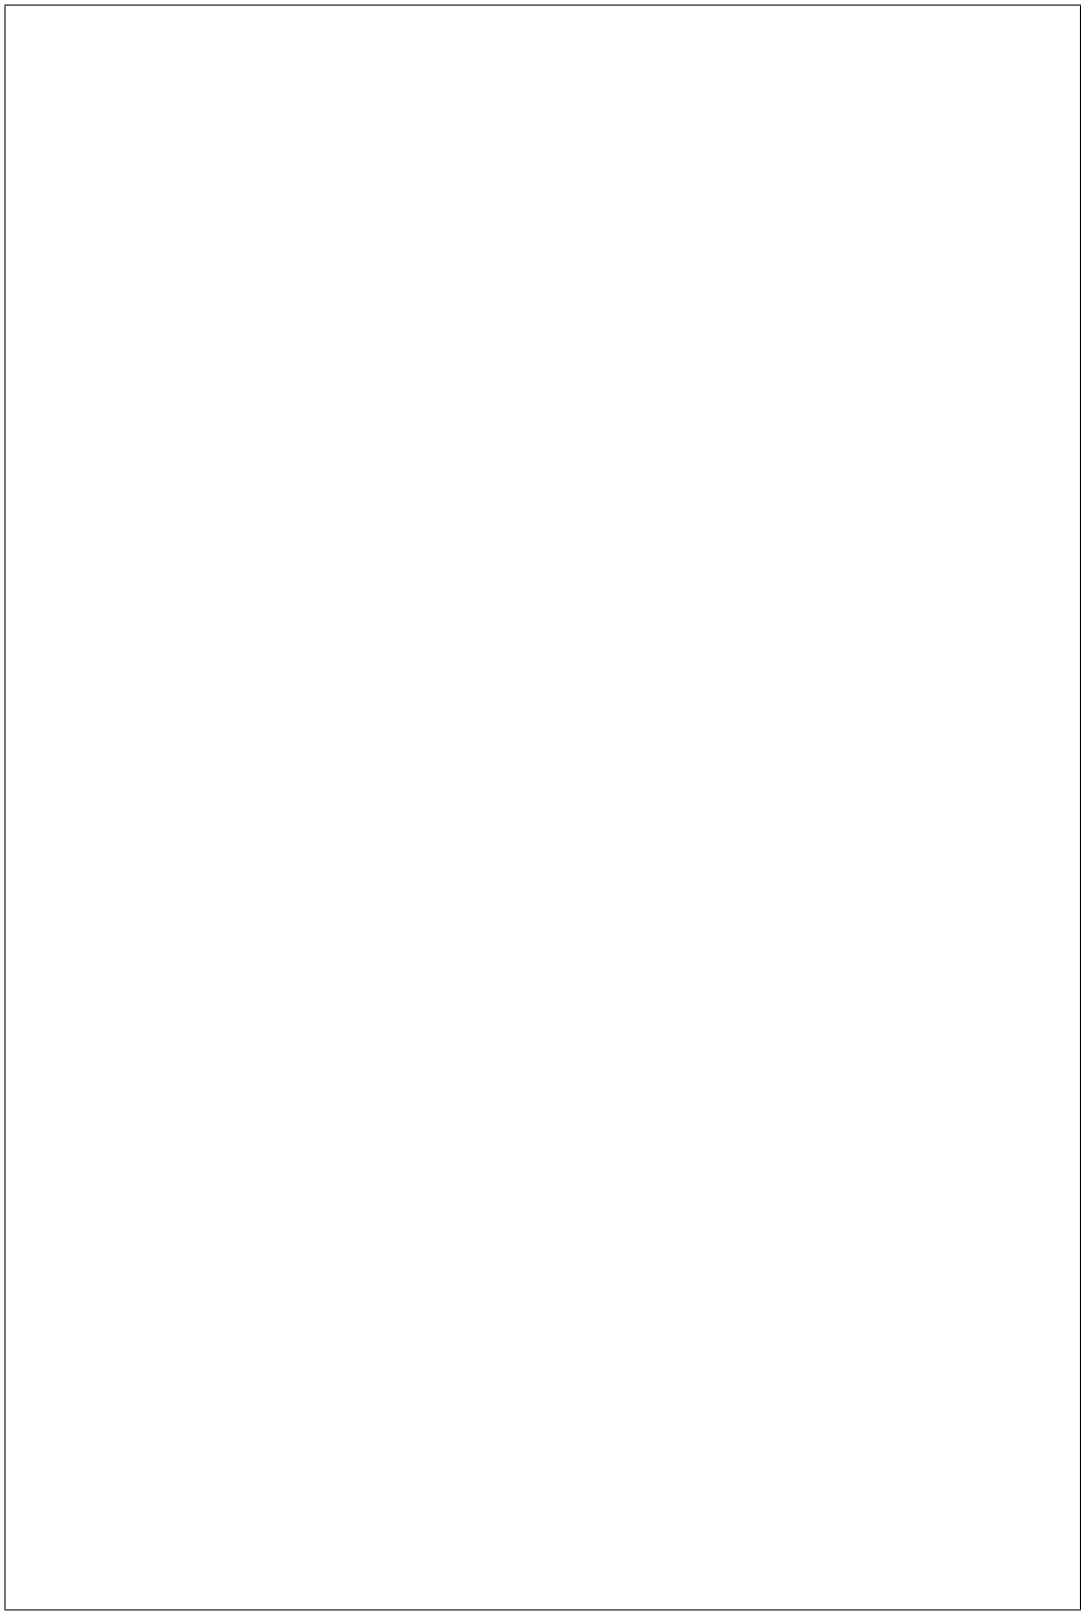

For $n = 2, 3$, we can solve the characteristic equation to get eigenvalues. For $n \geq 4$ there are better numerical methods.

31.4 Simple properties

For a square matrix A:

1. A and A^T have the same eigenvalues.

2. A is singular if and only if $\lambda = 0$ is an eigenvalue of A.

3. If λ is an eigenvalue of A, then λ^2 is an eigenvalue of A^2, and $1/\lambda$ is an eigenvalue of A^{-1} when A is non-singular.

4. If λ is an eigenvalue of A, then $\lambda - m$ is an eigenvalue of $A - mI$, for any scalar m.

31.5 Eigenvectors corresponding to distinct eigenvalues are linearly independent

If $\lambda_1, \lambda_2, \ldots, \lambda_k$ are distinct eigenvalues of A, with corresponding eigenvectors $\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_k$ (such that $\boldsymbol{v}_i$ corresponds to λ_i), then the set of eigenvectors $\{\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_k\}$ is linearly independent.

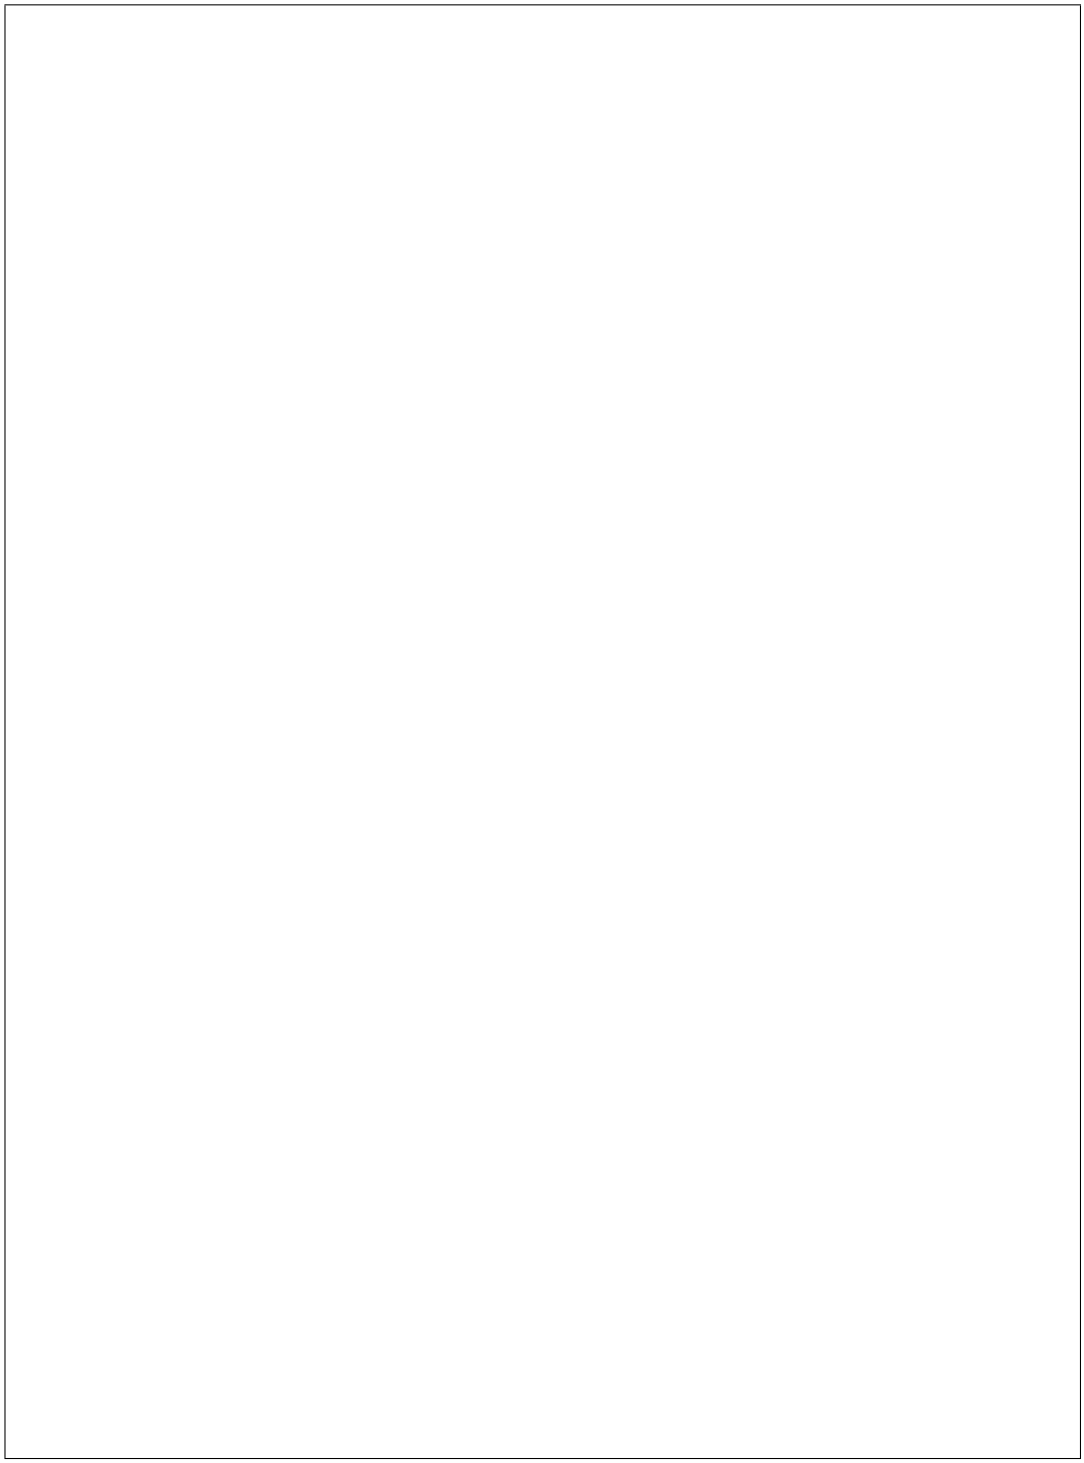

32 Diagonalisation

By the end of this section, you should be able to answer the following questions:

- How do you find a matrix P which diagonalises a given matrix A?
- How do you determine if A is diagonalisable?
- What are two applications of diagonalisation?

A square matrix A is *diagonalisable* if there is a non-singular matrix P such that $P^{-1}AP$ is a diagonal matrix. Here we consider the question: given a matrix, is it diagonalisable? If so, how do we find P?

The secret to constructing such a P is to let the columns of P be the eigenvectors of A. We immediately have that $AP = PD$, where D is a diagonal matrix with eigenvalues on the diagonal. We know from section 31.2 on page 198 that P is invertible if and only if the columns of P are linearly independent. Hence, we have the following result:

The $n \times n$ matrix A is diagonalisable if and only if A has n linearly independent eigenvectors.

Is the matrix $A = \begin{pmatrix} -3 & 1 & 0 \\ 1 & -2 & 1 \\ 0 & 1 & -3 \end{pmatrix}$ diagonalisable?

32.1 Similar matrices

Two matrices A and B are *similar* if there is a non-singular matrix P such that $B = P^{-1}AP$.

The statements "A is diagonalisable" and "A is similar to a diagonal matrix" are equivalent.

32.1.1 Theorem (similar matrices)

Similar matrices have the same eigenvalues.

In fact, if $B = P^{-1}AP$ and $\boldsymbol{v}$ is an eigenvector of A corresponding to eigenvalue λ, then $P^{-1}\boldsymbol{v}$ is an eigenvector of B corresponding to eigenvalue λ. This is because

$$\begin{aligned} B(P^{-1}\boldsymbol{v}) &= (P^{-1}AP)P^{-1}\boldsymbol{v} \\ &= P^{-1}(A\boldsymbol{v}) \\ &= P^{-1}(\lambda\boldsymbol{v}) \\ &= \lambda(P^{-1}\boldsymbol{v}) \end{aligned}$$

32.2 A closer look at the diagonal matrix

Let the matrix A be $n \times n$ with n linearly independent eigenvectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n$ corresponding to eigenvalues $\lambda_1, \ldots, \lambda_n$. Let

$$P = (\boldsymbol{v}_1 | \ldots | \boldsymbol{v}_n)$$

be the $n \times n$ matrix whose columns are the eigenvectors. Then

$$P^{-1}AP = \begin{pmatrix} \lambda_1 & 0 & \ldots & 0 \\ 0 & \lambda_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \lambda_n \end{pmatrix},$$

the diagonal matrix with the eigenvalues down the main diagonal. The important point here is the order in which the eigenvalues appear. They correspond to the order in which the associated eigenvectors appear in the columns of P.

32.3 Diagonalisability

We know that an $n \times n$ matrix A is diagonalisable if and only if A has n linearly independent eigenvectors.

Now say $\lambda_1, \ldots, \lambda_m$ are *distinct* eigenvalues of A, with corresponding eigenvectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m$. Then we have also seen that $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_m$ are linearly independent.

Hence if A is $n \times n$ with n distinct eigenvalues, then A is diagonalisable.

The question remains, if A has fewer than n distinct eigenvalues, how do we know if A is diagonalisable?

32.3.1 Example

Let $A = \begin{pmatrix} 2 & 1 & 3 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ and $B = \begin{pmatrix} 2 & 1 & 3 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}$.

Easy to see the characteristic equation of both A and B is $(2-\lambda)(1-\lambda)^2 = 0$, so $\lambda = 2, 1, 1$.

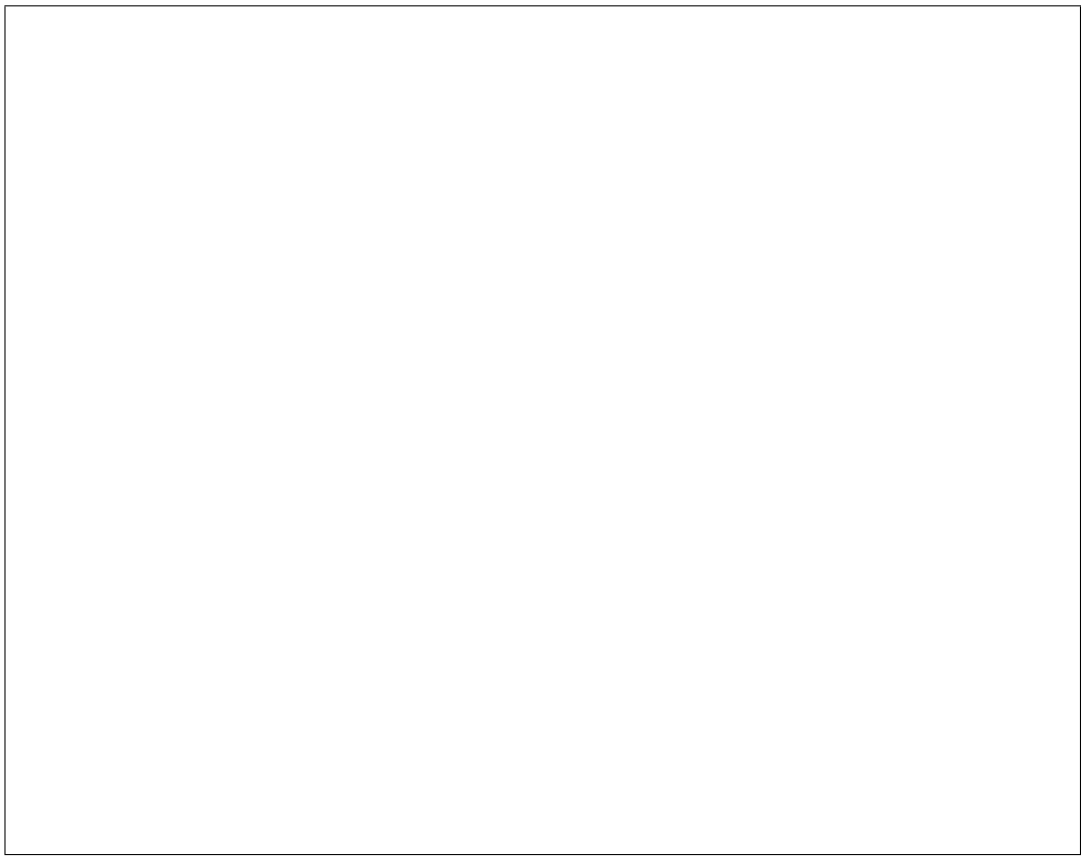

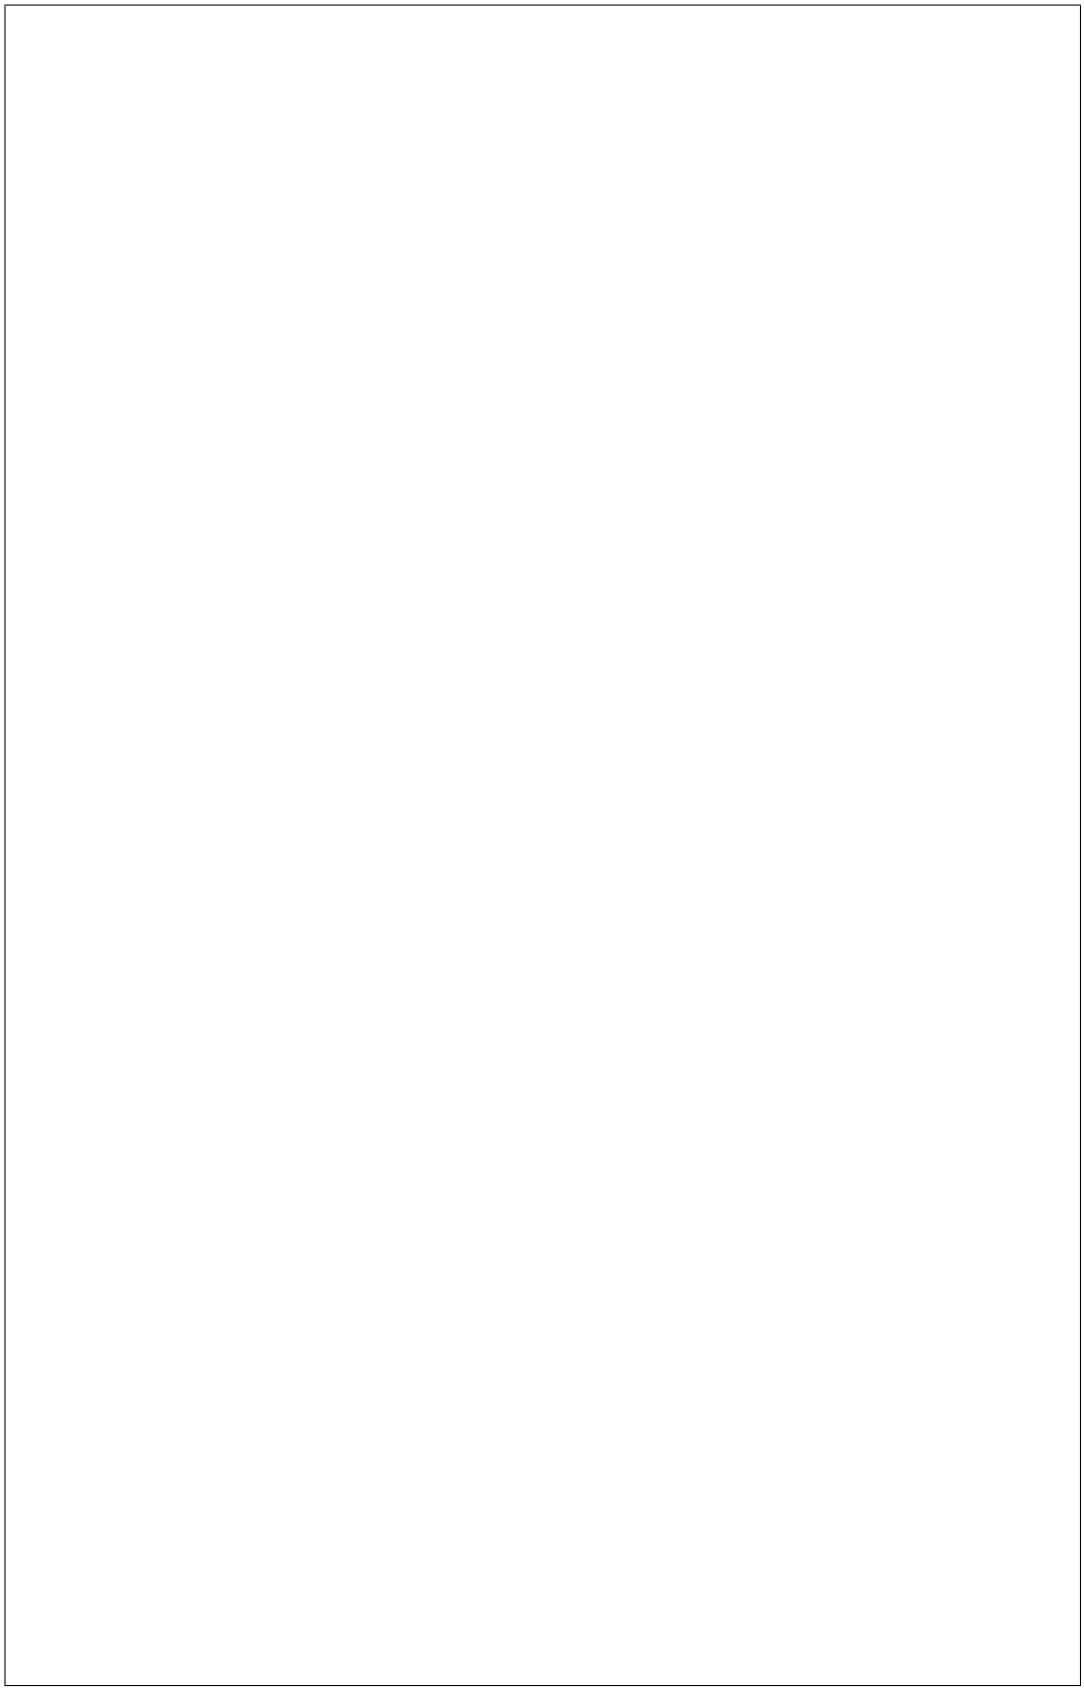

32.4 Algebraic and geometric multiplicity

If we are only interested in finding out whether or not a matrix is diagonalisable, then we need to know the dimension of each eigenspace. There is one theorem (which we will not prove!) that states:

If λ_i is an eigenvalue, then the dimension of the corresponding eigenspace cannot be greater than the number of times $(\lambda - \lambda_i)$ appears as a factor in the characteristic polynomial.

We often use the following terminology:

- The *geometric multiplicity* of the eigenvalue λ_i is the dimension of the eigenspace corresponding to λ_i.
- The *algebraic multiplicity* of the eigenvalue λ_i is the number of times $(\lambda - \lambda_i)$ appears as a factor in the characteristic polynomial.

The main result is the following:

A square matrix is diagonalisable if and only if the geometric and algebraic multiplicities are equal for every eigenvalue.

Note that the geometric multiplicity of λ_i is equal to $\text{nullity}(A - \lambda_i I)$. If A is $n \times n$, then the result at the bottom of page 197 tells us that

$$\text{nullity}(A - \lambda_i I) = n - \text{rank}(A - \lambda_i I).$$

In practice, we can determine the geometric multiplicity of λ_i by subtracting the number of non-zero rows in the r.e.f. of $(A - \lambda_i)$ from n. We then compare this number with the number of factors of $(\lambda - \lambda_i)$ to determine whether or not A is diagonalisable.

One of many corollaries to this result is that the geometric multiplicities of A and A^T are equal.

32.5 Applications of diagonalisability

32.5.1 Systems of differential equations

For a system of coupled differential equations which can be written in matrix form as

$$\dot{\boldsymbol{x}} = A\boldsymbol{x}$$

(where $\boldsymbol{x} = (x_1, \ldots, x_n)^T$, $\dot{\boldsymbol{x}} = (\dot{x}_1, \ldots, \dot{x}_n)^T$),
if A can be diagonalised, say $P^{-1}AP = D$ with D diagonal, then make the substitution $\boldsymbol{x} = P\boldsymbol{y}$. This yields

$$\dot{\boldsymbol{y}} = D\boldsymbol{y}$$

which is easily solved.

32.5.2 Matrix powers

If A is diagonalisable, say $P^{-1}AP = D$ with D diagonal, then

$$A^n = PD^nP^{-1}.$$

This gives an easy way to calculate A^n.

33 Orthogonal Diagonalisation

By the end of this section, you should be able to answer the following questions:

- What is a symmetric matrix?
- What is an orthogonal matrix?
- How do you diagonalise symmetric matrices?

Given an $n \times n$ matrix A, we call A *orthogonally diagonalisable* if there exists an orthogonal matrix P such that $P^{-1}AP = P^TAP$ is diagonal. To understand this, we first need to know what is meant by an orthogonal matrix.

33.1 Orthogonal matrices

An *orthogonal* matrix is a real square matrix Q such that the columns of Q are mutually orthogonal unit vectors (i.e. $\boldsymbol{v}_i \cdot \boldsymbol{v}_j = 0$ if $i \neq j$, and $|\boldsymbol{v}_i| = 1$).

Note that mutually orthogonal unit vectors are called *orthonormal*.

An orthogonal matrix is then a real square matrix Q such that $Q^{-1} = Q^T$. Note also that $\det(Q) = \pm 1$.

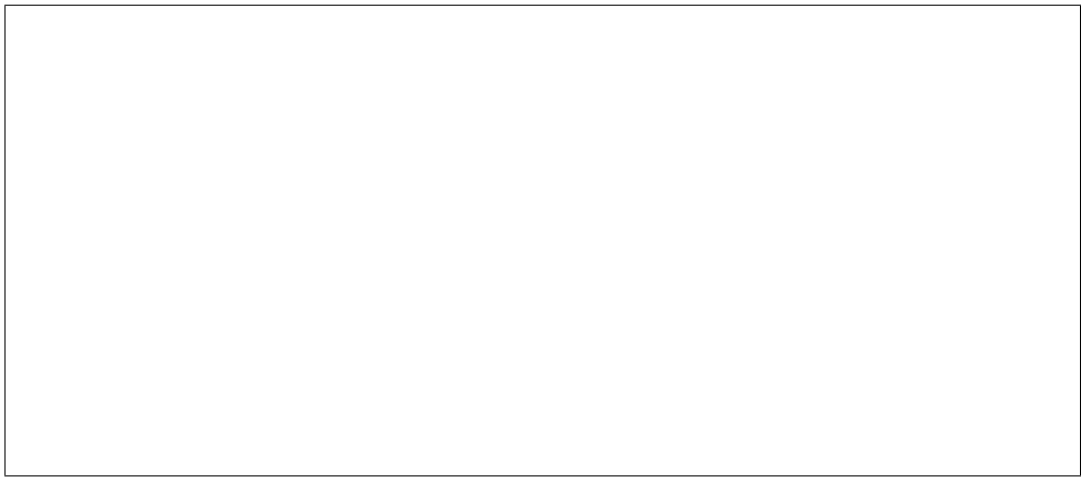

33.2 Symmetric matrices

A matrix A is *symmetric* if and only if $A = A^T$. Symmetric matrices are easy to identify due to their "mirror symmetry" about the main diagonal. For example, we can tell by inspection that $A = \begin{pmatrix} -3 & 1 & 0 \\ 1 & -2 & 1 \\ 0 & 1 & -3 \end{pmatrix}$ is symmetric.

33.2.1 If A is real symmetric, then the eigenvectors corresponding to different eigenvalues are orthogonal.

Proof:

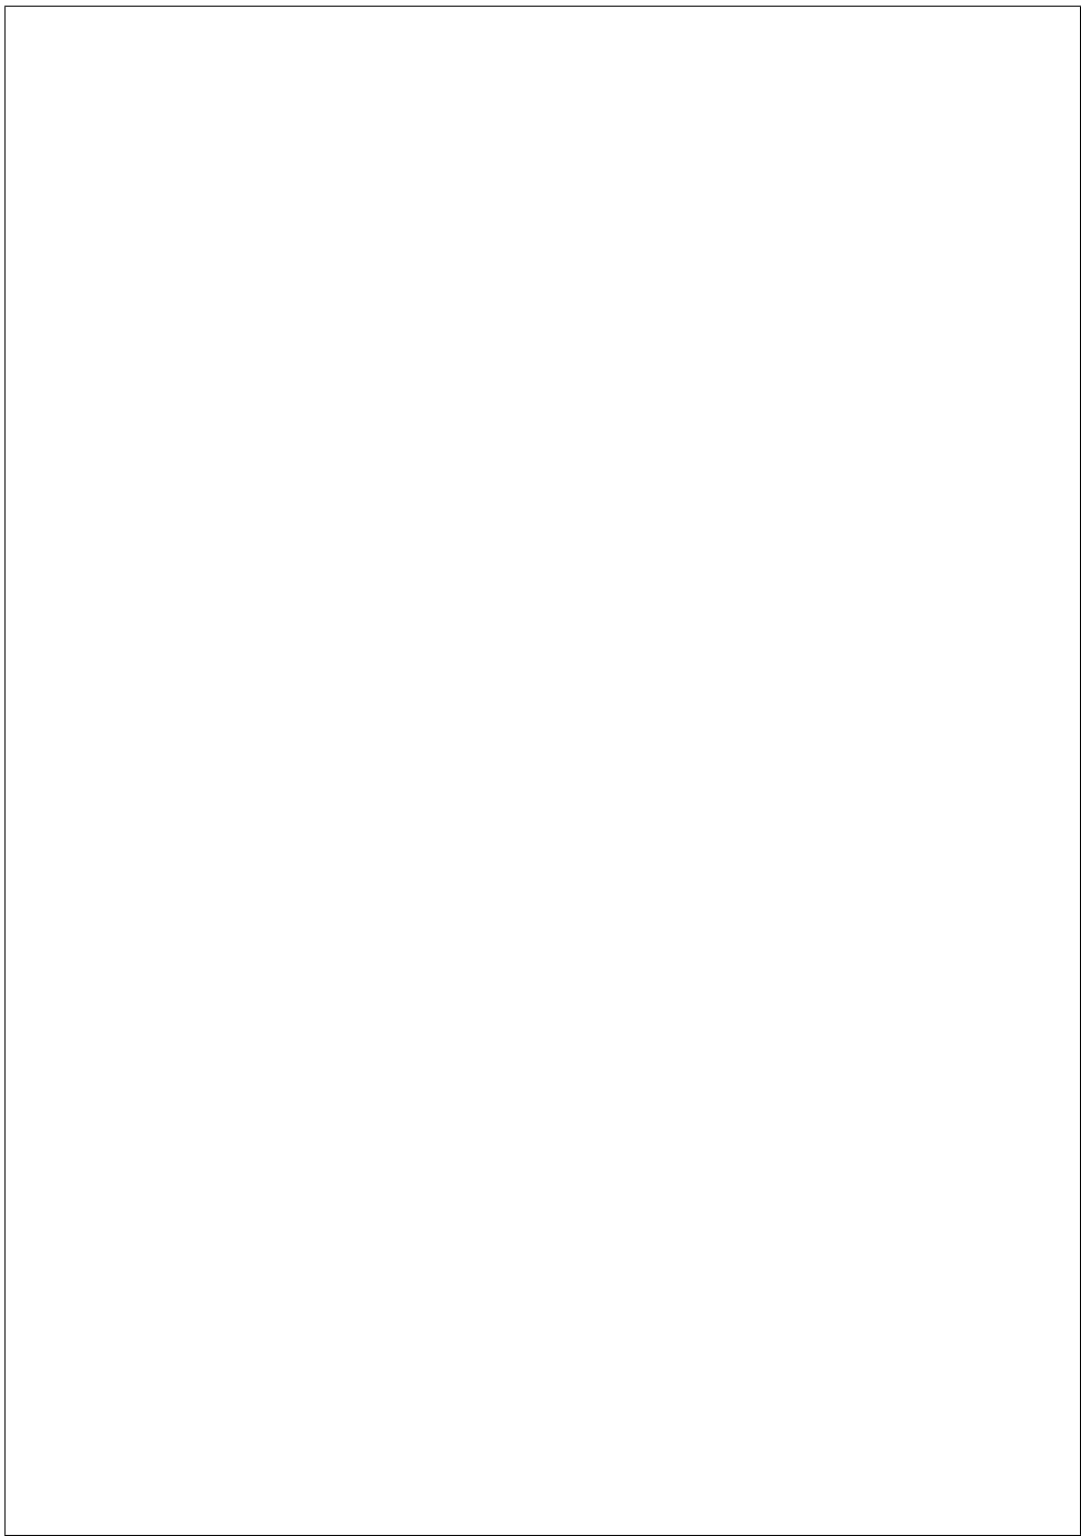

33.2.2 Real symmetric matrices are orthogonally diagonalisable

It is straightforward to show that if a matrix is orthogonally diagonalisable, then it is symmetric:

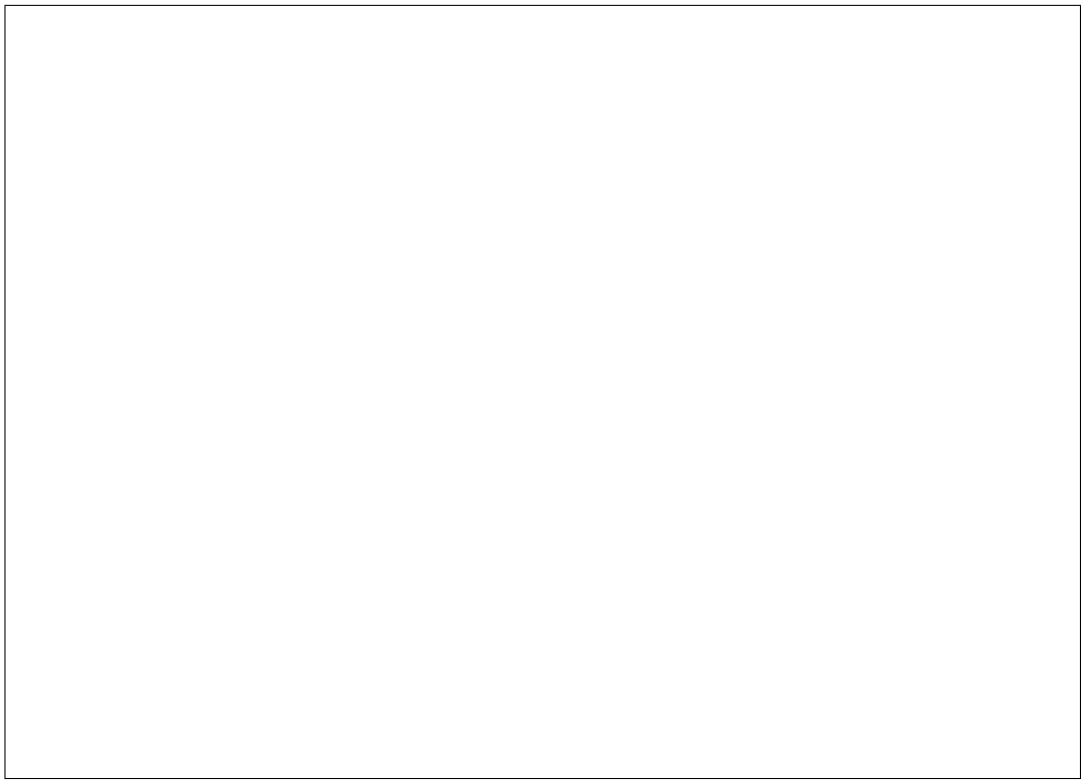

In fact, the converse is also true (although difficult to prove), giving us the amazing result:

An $n \times n$ real matrix is orthogonally diagonalisable if and only if it symmetric.

The significance of this is that a symmetric matrix is *always* diagonalisable by an orthogonal matrix.

33.2.3 Eigenvectors and eigenvalues

Here we state two results about any symmetric matrix A without proof:

(1) All the eigenvalues of A are real;

(2) A has n linearly independent eigenvectors.

33.2.4 Example

Let $A = \begin{pmatrix} -3 & 1 & 0 \\ 1 & -2 & 1 \\ 0 & 1 & -3 \end{pmatrix}$ (see previous examples).

We already know the eigenvalues are $-3, -1, -4$ with corresponding eigenvectors

$$\boldsymbol{v}_1 = \begin{pmatrix} 1 \\ 0 \\ -1 \end{pmatrix}, \quad \boldsymbol{v}_2 = \begin{pmatrix} 1 \\ 2 \\ 1 \end{pmatrix}, \quad \boldsymbol{v}_3 = \begin{pmatrix} 1 \\ -1 \\ 1 \end{pmatrix}.$$

Note that A is real symmetric, so $\boldsymbol{v}_1$, $\boldsymbol{v}_2$ and $\boldsymbol{v}_3$ should be pairwise orthogonal.

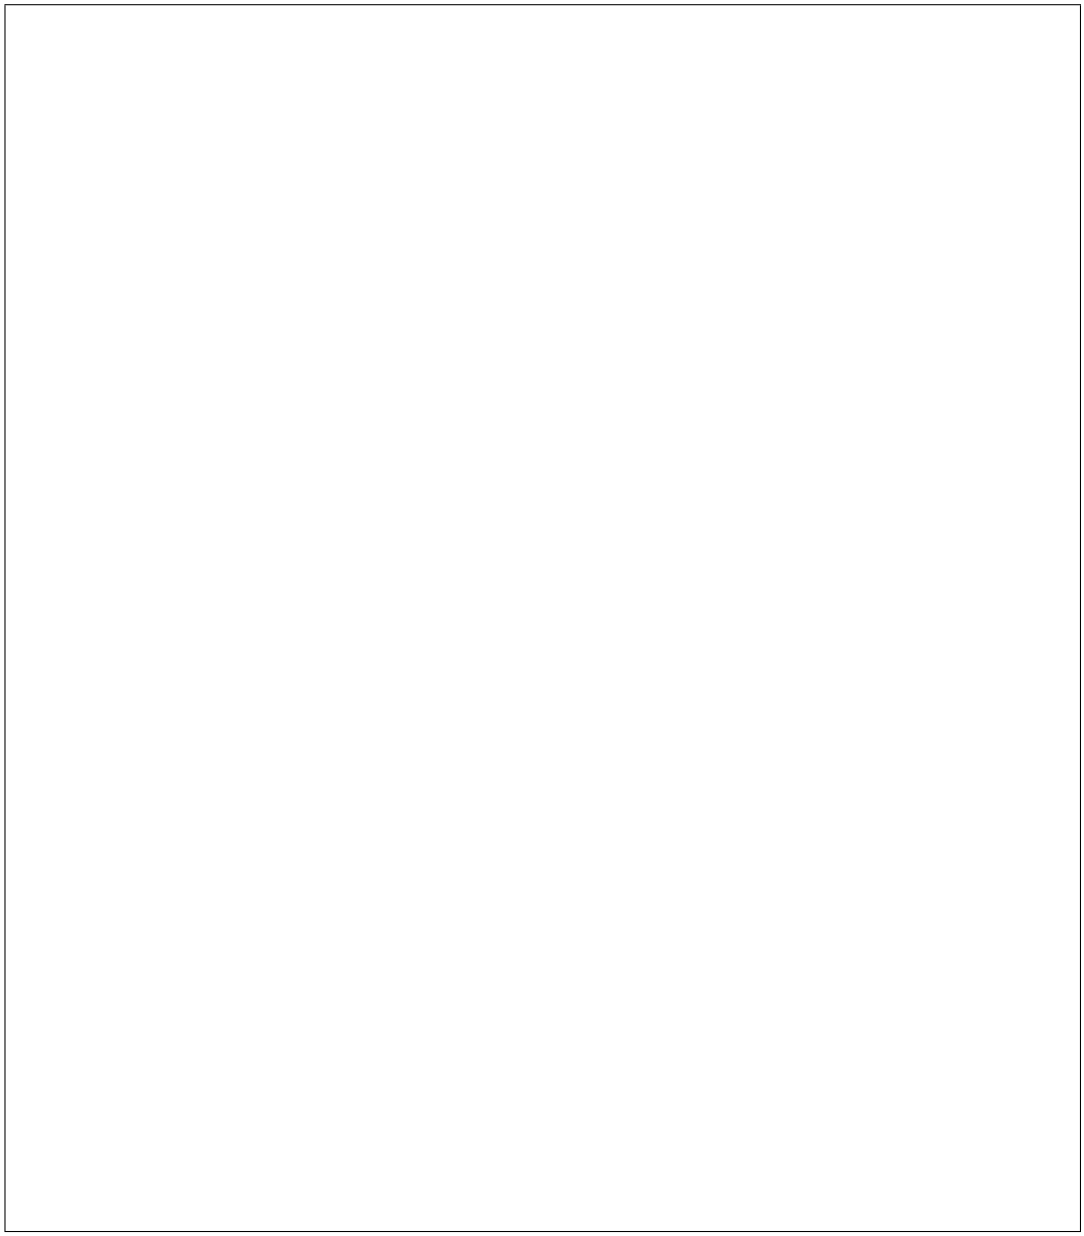

34 Quadratic forms

By the end of this section, you should be able to answer the following questions:

- What is a quadratic form?
- How do you diagonalise quadratic forms?
- How can you use diagonalisation of two variable quadratic forms to identify conic sections?
- What are quadric surfaces?

This section presents a novel application of orthogonal diagonalisation as a way of identifying conic sections. We also mention the generalisation to three dimensions and how, in principle, we could identify quadric surfaces, although the details in this case can become quite messy.

The majority of this section is based on the section on quadratic forms in the MATH2000 recommended text "Elementary Linear Algebra (Applications Version)" by Anton and Rorres, pages 479–502.

34.1 Definition

Consider n real variables $x_1, x_2, \ldots, x_n$. A function of the form $\sum_{i=1}^{n}\sum_{j=1}^{n} a_{ij}x_i x_j$ is called a *quadratic form*, where the a_{ij} are real constants.

For example, the most general quadratic form in the variables x and y is

$$Q(x, y) = ax^2 + by^2 + cxy.$$

In the three variables x, y and z, the most general quadratic form is

$$Q(x, y, z) = ax^2 + by^2 + cz^2 + dxy + exz + fyz,$$

where in both cases a, b, c, d, e, f are all constants. It is possible to express quadratic forms in n variables as a matrix product $\boldsymbol{v}^T A\boldsymbol{v}$, where $\boldsymbol{v}$ is a vector with the n variables as entries and A is a symmetric matrix.

The two variable quadratic form above can be expressed as

$$Q(x, y) = \begin{pmatrix} x & y \end{pmatrix} \begin{pmatrix} a & c/2 \\ c/2 & b \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix},$$

while the three variable quadratic form given above can be written as

$$Q(x,y,z) = \begin{pmatrix} x & y & z \end{pmatrix} \begin{pmatrix} a & d/2 & e/2 \\ d/2 & b & f/2 \\ e/2 & f/2 & c \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix}.$$

As an exercise, trying verifying this by expanding out both expressions. Observe that in both cases the diagonal entries of the matrix are the coefficients of the square terms and the off-diagonal entries in the matrix are the coefficients of the cross-terms.

34.1.1 Give the matrix representation of the quadratic form $2x^2 + 6xy - 7y^2$.

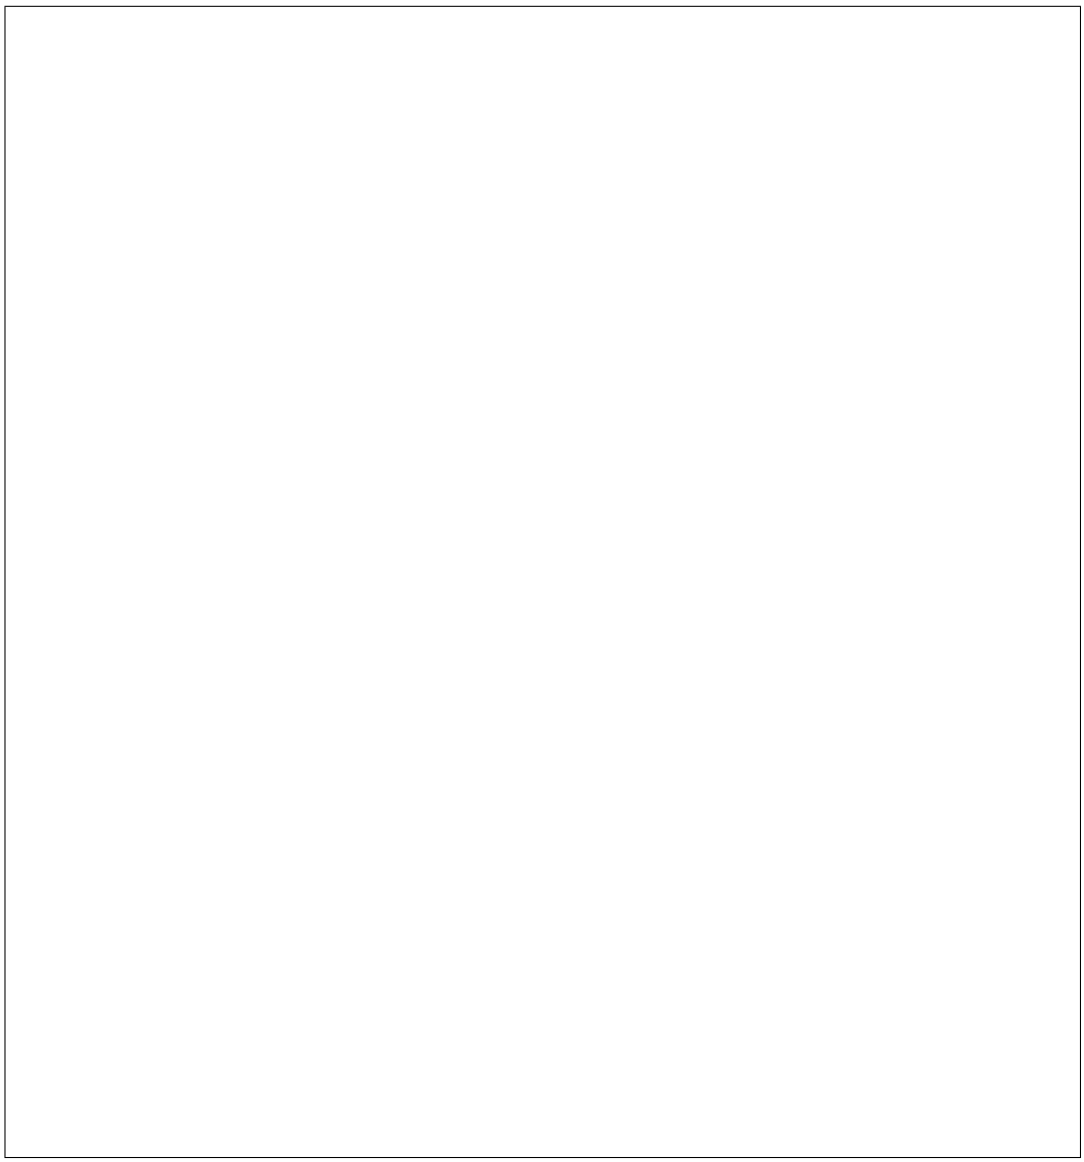

34.2 Diagonalising quadratic forms

Since we know we can always orthogonally diagonalise a symmetric matrix, if we do this to the symmetric matrix in the matrix representation of the quadratic form, we can reduce the quadratic form to a sum of square terms.

We shall demonstrate this by example:

34.2.1 Express $-3x^2 - 2y^2 - 3z^2 + 2xy + 2yz$ exclusively as a sum of square terms.

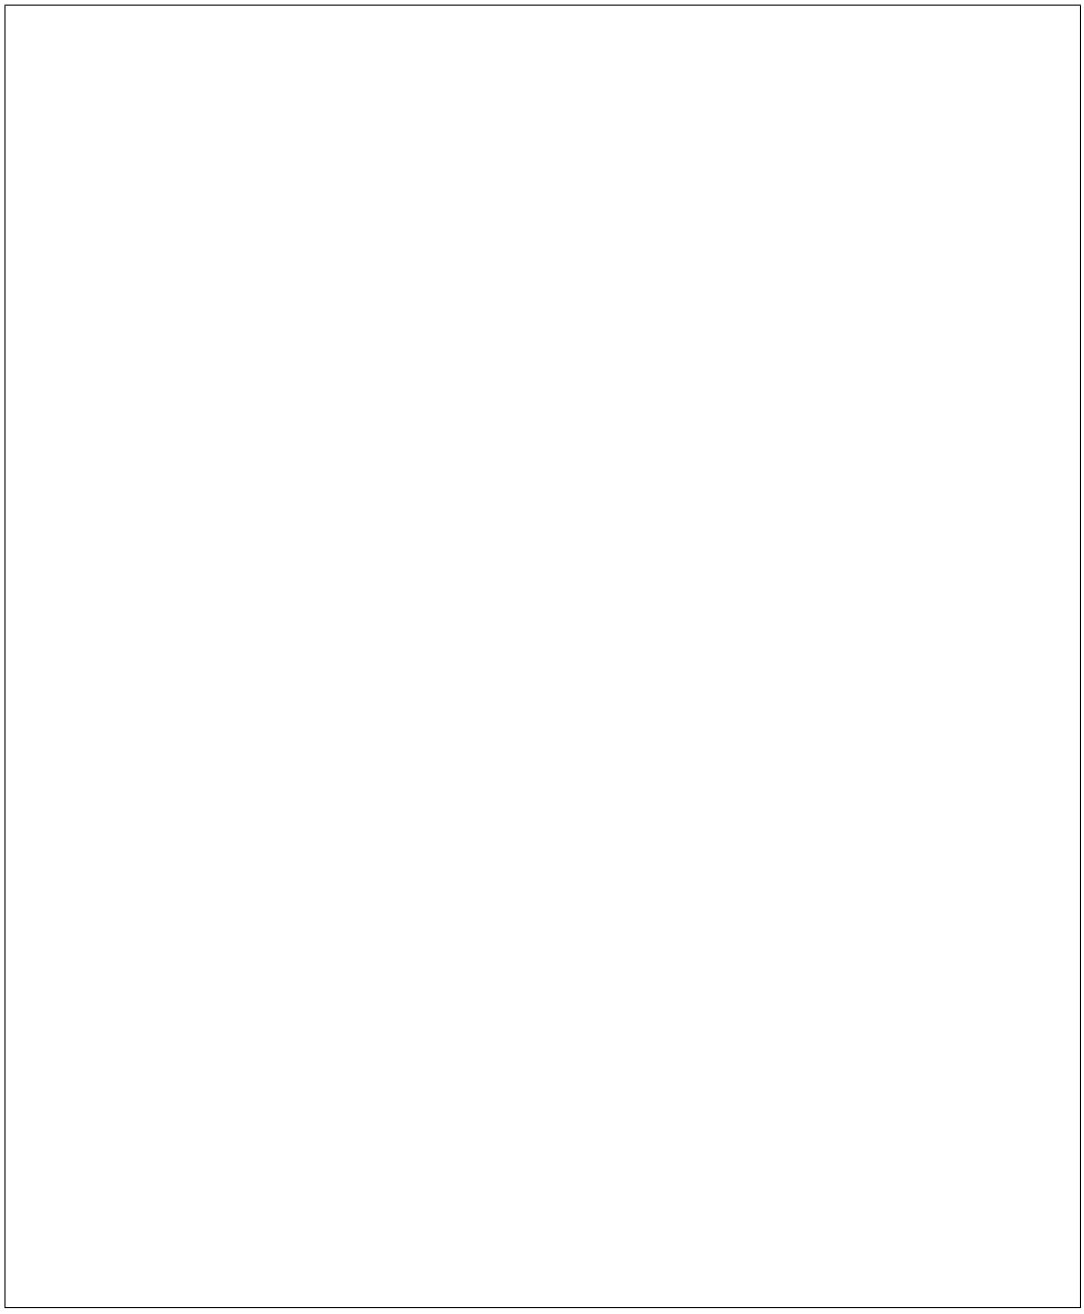

34.3 Quadratic equations and conic sections

We now restrict our attention to two dimensions, by investigating quadratic equations, which are equations of the form

$$ax^2 + by^2 + cxy + dx + ey + f = 0,$$

where $a, b, c, d, e, f \in \mathbb{R}$.

Graphs of quadratic equations are known as *conic sections*, because they can be realised as the intersection of a plane and a double cone in three dimensions. The most interesting of these are the so-called *non-degenerate* conic sections[2]. A non-degenerate conic section is in standard position relative to the coordinate axes if its equation can be expressed in one of the following forms:

- $\dfrac{x^2}{k^2} + \dfrac{y^2}{l^2} = 1; k, l > 0,$
- $\dfrac{x^2}{k^2} - \dfrac{y^2}{l^2} = 1$ or $\dfrac{y^2}{l^2} - \dfrac{x^2}{k^2} = 1; k, l > 0,$
- $x^2 = ky$ or $y^2 = kx; k \neq 0.$

The key observation here is that conic sections in standard form have no cross-terms. Given a quadratic equation with cross-terms in the associated quadratic form, we can *change variables* to remove the cross-terms by orthogonal diagonalisation. Due to the defining property of rotation matrices, an orthogonal matrix P always corresponds to a rotation, provided $\det(P) = 1$ (not -1). Hence, we have the following.

Changing variables by orthogonal diagonalisation corresponds to a rotation of the coordinate axes. If P is the orthogonal (rotation) matrix, then the new coordinates (u, v) can be expressed in terms of the old coordinates (x, y) as

$$\begin{pmatrix} u \\ v \end{pmatrix} = P^T \begin{pmatrix} x \\ y \end{pmatrix}.$$

Another important observation is that there is never an occurance of x^2 and x in the standard form (or y^2 and y). As a general rule, given a quadratic equation (even after changing variables from orthogonal diagonalisation), if we have terms such as x^2 and x (or similar terms involving new variables) we can *complete the square* to be left with only a square term. We have the following.

Completing the square in a quadratic equation corresponds to translating (or shifting) the coordinate axes.

[2]There are also *degenerate* (points, lines) and *imaginary* (without real graphs) conic sections.

In summary, to identify a quadratic equation as a conic section, we follow these steps:

1. Write the quadratic equation
$$ax^2 + by^2 + cxy + dx + ey + f = 0$$
in the matrix form $\boldsymbol{x}^T A\boldsymbol{x} + K\boldsymbol{x} + f = 0$, where $\boldsymbol{x} = \begin{pmatrix} x \\ y \end{pmatrix}$ and $K = \begin{pmatrix} d & e \end{pmatrix}$.

2. Find a matrix P that orthogonally diagonalises A, so $A = PDP^T$. You may need to swap columns of P to ensure that $\det(P) = 1$ (and hence corresponds to a rotation).

3. Define new variables u, v such that $\boldsymbol{v} = \begin{pmatrix} u \\ v \end{pmatrix} = P^T\boldsymbol{x} \Rightarrow \boldsymbol{x} = P\boldsymbol{v}$.

4. Substitute $\boldsymbol{v}$ into the matrix form of the equation, giving
$$\boldsymbol{v}^T D\boldsymbol{v} + KP\boldsymbol{v} + f = 0.$$

5. Complete the square if required. This is necessary if u^2 and u are both present (or v^2 and v). This defines a new set of variables s, t by translating u, v. The translations will be of the form $s = \alpha u + \beta$, $t = \gamma v + \delta$.

6. If it is a non-degenerate conic, the final equation in s and t should be a conic section in standard form.

34.3.1 Describe the conic whose equation is $x^2 + y^2 + 2xy - 3x - 5y + 4 = 0$.

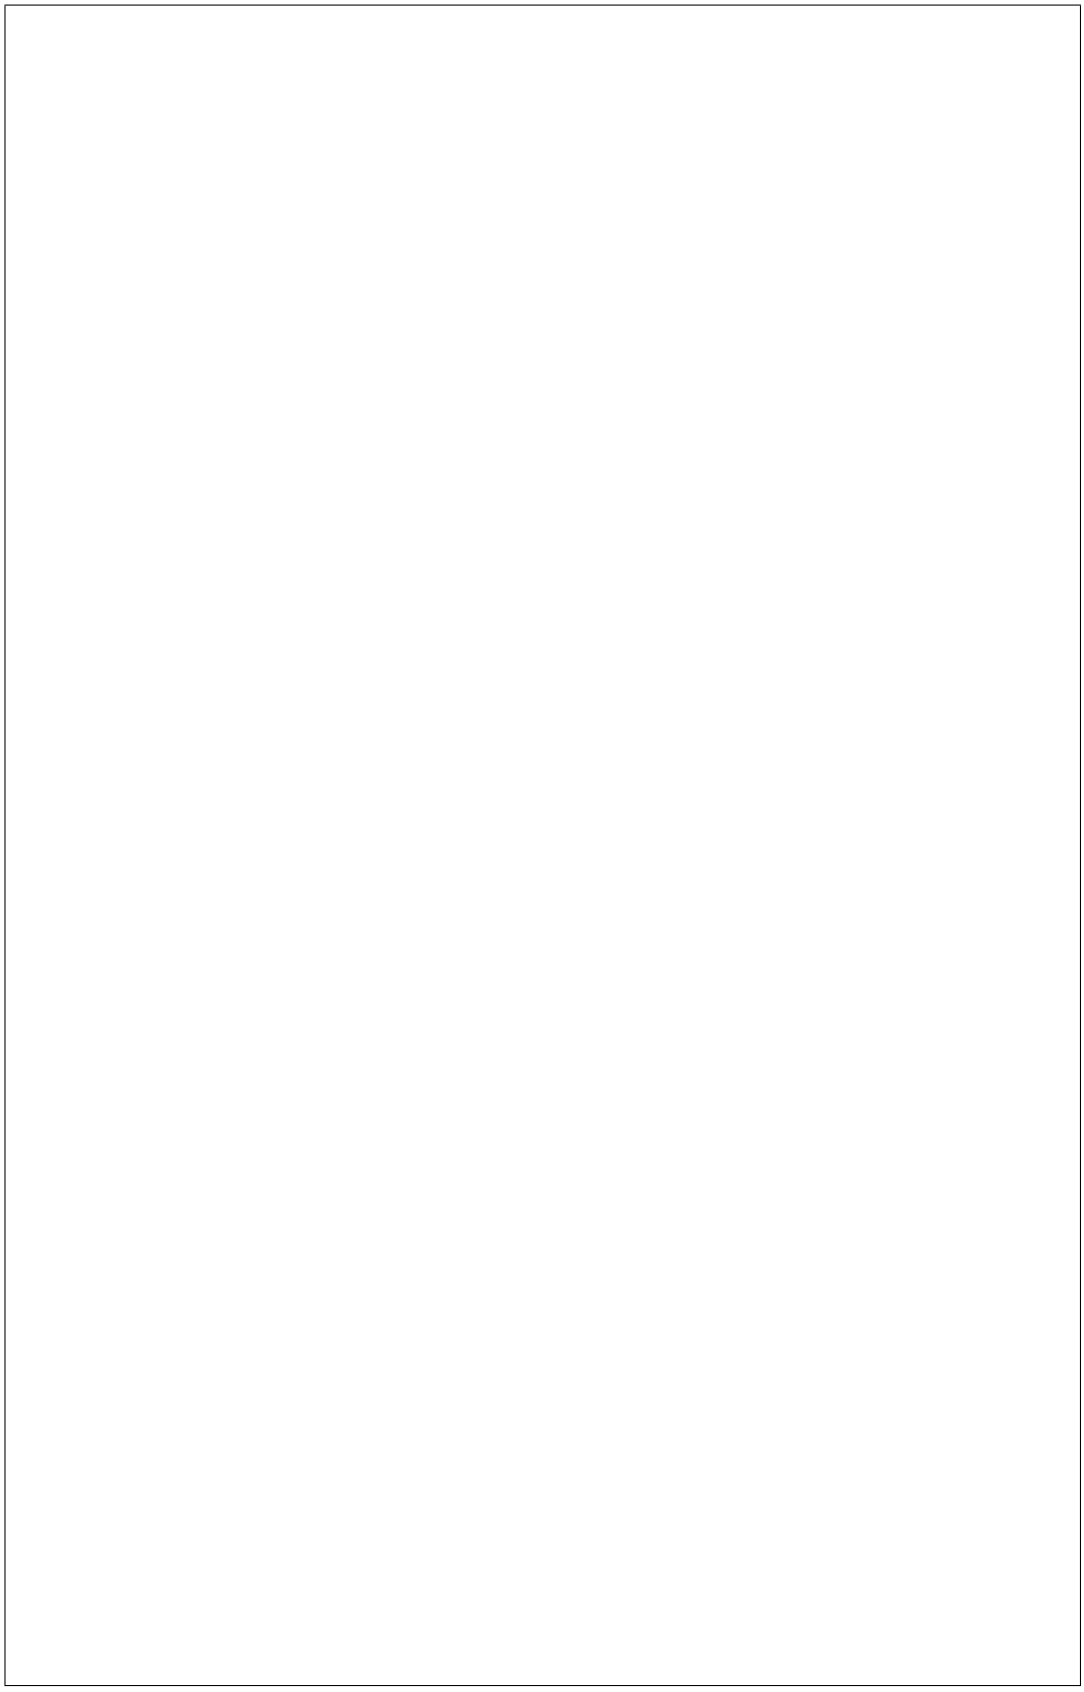

34.4 Quadric surfaces

It turns out we can similarly use orthogonal diagonalisation of 3×3 matrices to simplify and ultimately identify surfaces whose general equation is of the form

$$ax^2 + by^2 + cz^2 + 2dxy + 2exz + 2fyz + gx + hy + iz + j = 0,$$

where a, b, c, d, e, f are never all zero. Note that we can rewrite the equation in matrix form as

$$\begin{pmatrix} x & y & z \end{pmatrix} \begin{pmatrix} a & d & e \\ d & b & f \\ e & f & c \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} g & h & i \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} + j = 0.$$

If we then orthogonally diagonalise the 3×3 matrix, then complete any squares that are left over, we end up being able to identify the surface as one of the following forms:

- $\dfrac{x^2}{l^2} + \dfrac{y^2}{m^2} + \dfrac{z^2}{n^2} = 1,$
- $z^2 = \dfrac{x^2}{l^2} + \dfrac{y^2}{m^2},$
- $\dfrac{x^2}{l^2} + \dfrac{y^2}{m^2} - \dfrac{z^2}{n^2} = 1,$
- $z = \dfrac{x^2}{l^2} + \dfrac{y^2}{m^2},$
- $\dfrac{z^2}{l^2} - \dfrac{x^2}{m^2} - \dfrac{y^2}{n^2} = 1,$
- $z = \dfrac{y^2}{m^2} - \dfrac{x^2}{l^2}.$

As in the two dimensional case, the orthogonal diagonalisation has the effect of rotating the axes, provided the orthogonal matrix P has $\det P = 1$, which we can choose by carefully ordering the columns. Completing the square has the effect of shifting the axes.

You should be aware that these techniques are available in order to simplify and identify algebraic expressions representing surfaces. The three dimensional case can often become quite complicated. You will not be expected to identify quadric surfaces in an exam.

35 Power method

By the end of this section, you should be able to answer the following questions:

- What is the power method and what does it do?
- Under what conditions can it fail?
- What is deflation, and how does it work in conjunction with the power method?

In applications we sometimes need to find eigenvalues and eigenvectors of a large square matrix. In these cases it is usually impractical, or more to the point not computationally feasible, to find the roots of the characteristic polynomial. Instead, we are forced to rely on computational techniques which estimate eigenvalues and eigenvectors. The power method is one such technique which estimates the largest eigenvalue (provided it is unique) and its corresponding eigenvector.

35.1 Dominant eigenvalue

Let A be an $n \times n$ matrix with eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_n$ such that

$$|\lambda_1| > |\lambda_2| \geq \ldots \geq |\lambda_n|.$$

The eigenvalue λ_1 of the matrix A is called the *dominant eigenvalue* of A. The eigenvector $\boldsymbol{v}_1$ corresponding to λ_1 is called the *dominant eigenvector*.

35.1.1 Example

Identify the dominant eigenvalue and eigenvector of the matrix $\begin{pmatrix} -3 & 1 & 0 \\ 1 & -2 & 1 \\ 0 & 1 & -3 \end{pmatrix}$.

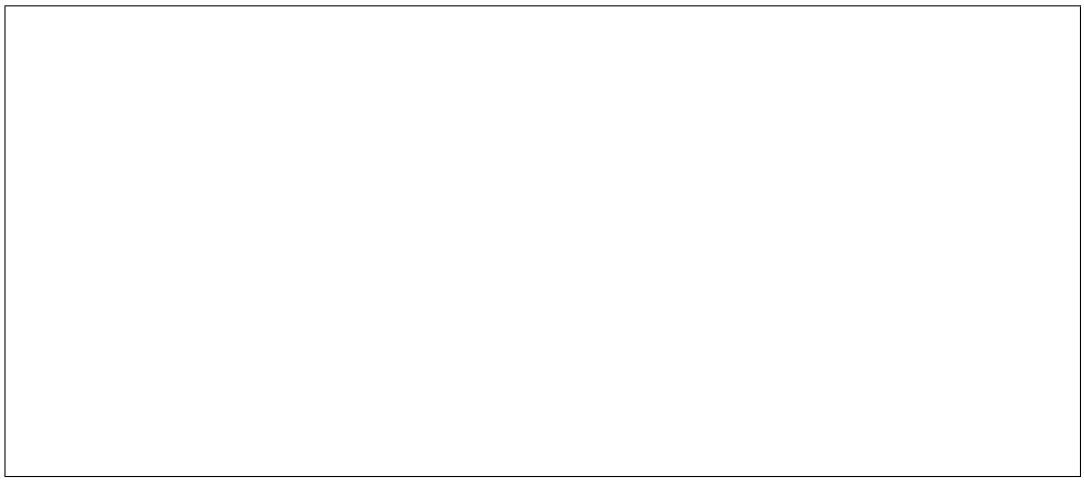

35.2 The algorithm

Form a sequence of vectors $\boldsymbol{u}_0, \boldsymbol{u}_1, \ldots, \boldsymbol{u}_k, \ldots$ where $\boldsymbol{u}_0$ is an (almost!) arbitrarily chosen vector, $\boldsymbol{u}_{k+1} = A\boldsymbol{u}_k$ (for $k \geq 0$). Then (usually) for k large,

(i) The dominant eigenvalue is $\lambda_1 \approx \dfrac{(\boldsymbol{u}_{k+1})_j}{(\boldsymbol{u}_k)_j}$, any $j \leq n$ with $(\boldsymbol{u}_k)_j \neq 0$ (usually we choose j so that $|(\boldsymbol{u}_{k+1})_j|$ is the largest possible),

(ii) $\boldsymbol{u}_k \approx$ dominant eigenvector.

35.2.1 Example

For $A = \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix}$, find the exact value of the dominant eigenvalue and eigenvector, then apply the power method approximation.

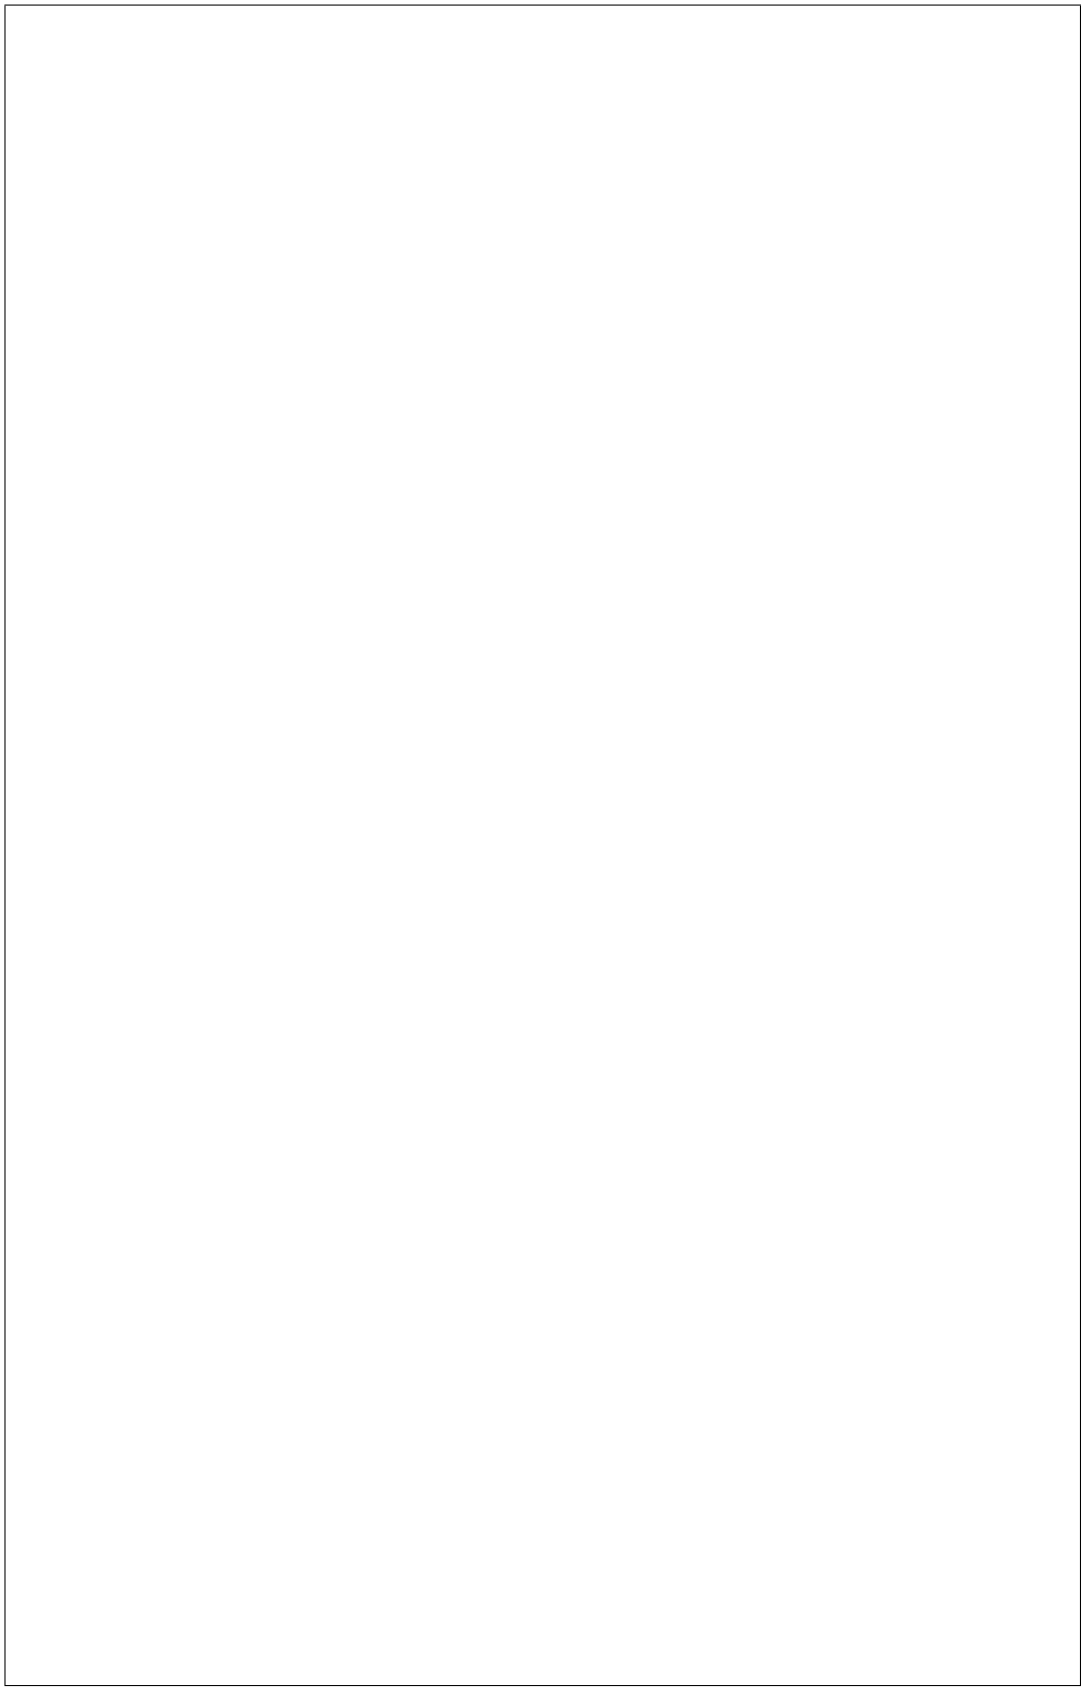

35.2.2 Assumptions

The power method depends on several assumptions:

1. There is a dominant eigenvalue.
2. The eigenvectors $\boldsymbol{v}_1, \boldsymbol{v}_2, \ldots, \boldsymbol{v}_n$ are linearly independent and hence form a basis for $\mathbb{R}^n$.
3. The chosen vector $\boldsymbol{u}_0$ that starts the iteration is non-zero and when written as a linear combination of the basis of eigenvectors, has a non-zero component of the dominant eigenvector.

35.2.3 Understanding the power method

Suppose λ_1 is the dominant eigenvalue of an $n \times n$ matrix A, so that

$$|\lambda_1| > |\lambda_2|, \ldots, |\lambda_n|$$

and hence $\lambda_1 \neq 0$. For simplicity, suppose that A has n linearly independent eigenvectors $\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n \in \mathbb{R}^n$. With n linearly independent vectors in $\mathbb{R}^n$, we have a basis, so any vector $\boldsymbol{u} \in \mathbb{R}^n$ can be written as a linear combination of the vectors in the basis. In particular, set

$$\boldsymbol{u}_0 = t_1\boldsymbol{v}_1 + t_2\boldsymbol{v}_2 + \cdots + t_n\boldsymbol{v}_n$$

for some scalars $t_1, \ldots, t_n$. Suppose $t_1 \neq 0$ (this turns out to be crucial). Then

$$\begin{aligned}
\boldsymbol{u}_1 &= A\boldsymbol{u}_0 = t_1A\boldsymbol{v}_1 + t_2A\boldsymbol{v}_2 + \cdots + t_nA\boldsymbol{v}_n = t_1\lambda_1\boldsymbol{v}_1 + t_2\lambda_2\boldsymbol{v}_2 + \cdots + t_n\lambda_n\boldsymbol{v}_n, \\
\boldsymbol{u}_2 &= A\boldsymbol{u}_1 = t_1\lambda_1A\boldsymbol{v}_1 + \cdots + t_n\lambda_nA\boldsymbol{v}_n = t_1\lambda_1^2\boldsymbol{v}_1 + \cdots + t_n\lambda_n^2\boldsymbol{v}_n
\end{aligned}$$

and in general

$$\begin{aligned}
\boldsymbol{u}_k &= t_1\lambda_1^k\boldsymbol{v}_1 + t_2\lambda_2^k\boldsymbol{v}_2 + \cdots + t_n\lambda_n^k\boldsymbol{v}_n \\
&= \lambda_1^k\left[t_1\boldsymbol{v}_1 + t_2\left(\frac{\lambda_2}{\lambda_1}\right)^k\boldsymbol{v}_2 + \cdots + t_n\left(\frac{\lambda_n}{\lambda_1}\right)^k\boldsymbol{v}_n\right].
\end{aligned}$$

Since $|\lambda_1| > |\lambda_2|, \ldots, |\lambda_n|$, $\left|\frac{\lambda_2}{\lambda_1}\right| < 1, \ldots, \left|\frac{\lambda_n}{\lambda_1}\right| < 1$ so that $\left(\frac{\lambda_i}{\lambda_1}\right)^k \to 0$ as $k \to \infty$. So for large k,

$$\boldsymbol{u}_k \approx \lambda_1^k t_1\boldsymbol{v}_1$$

i.e. $\boldsymbol{u}_k \approx$ eigenvector corresponding to λ_1.

Also, $\boldsymbol{u}_{k+1} \approx \lambda_1^{k+1} t_1 \boldsymbol{v}_1$, so

$$\frac{(\boldsymbol{u}_{k+1})_j}{(\boldsymbol{u}_k)_j} \approx \frac{(\lambda_1^{k+1} t_1 \boldsymbol{v}_1)_j}{(\lambda_1^k t_1 \boldsymbol{v}_1)_j} = \frac{\lambda_1^{k+1} t_1 (\boldsymbol{v}_1)_j}{\lambda_1^k t_1 (\boldsymbol{v}_1)_j} = \lambda_1.$$

Note this does not work if $t_1 = 0$, i.e. if $\boldsymbol{u}_0$ is a linear combination of only non-dominant eigenvectors.

35.3 Deflation

The power method gives only the dominant eigenvalue. For symmetric matrices, we can find the next most dominant one by *deflation*, based on the following.

If A is $n \times n$ with eigenvalues $\lambda_1, \lambda_2, \ldots, \lambda_n$ and $\boldsymbol{v}_1$ is an eigenvector corresponding to λ_1, then set

$$B = A - \left(\frac{\lambda_1}{\boldsymbol{v}_1^T \boldsymbol{v}_1} \right) \boldsymbol{v}_1 \boldsymbol{v}_1^T.$$

Note that $\boldsymbol{v}_1 \boldsymbol{v}_1^T$ is a symmetric $n \times n$ matrix, and hence B is symmetric.

If A is symmetric and $\boldsymbol{v}_i$ is an eigenvector of A corresponding to $\lambda_i \neq 0$, then $\boldsymbol{v}_i$ is also an eigenvector of B corresponding to λ_i.

The eigenvalues of B are $0, \lambda_2, \ldots, \lambda_n$.

For symmetric A with eigenvalues $\lambda_1, \ldots, \lambda_n$ where $|\lambda_1| > |\lambda_2| > \cdots > |\lambda_n|$, having used the power method to find an approx'n to λ_1 and $\boldsymbol{v}_1$, form

$$B = A - \left(\frac{\lambda_1}{\boldsymbol{v}_1^T \boldsymbol{v}_1} \right) \boldsymbol{v}_1 \boldsymbol{v}_1^T$$

and repeat the power method on B to find an approximation for λ_2 and $\boldsymbol{v}_2$.

In theory, you could repeat the power-deflation combination for other eigenvalues, but because the power method only approximates the dominant eigenvalue, we will be introducing some error into the method of deflation. Each time we repeat the process, the error not only propagates, but grows substantially.

35.3.1 Example

Apply deflation to the previous example of $A = \begin{pmatrix} 3 & 1 \\ 1 & 3 \end{pmatrix}$, then use the power method on the new matrix to approximate the next most dominant eigenvalue and corresponding eigenvector.

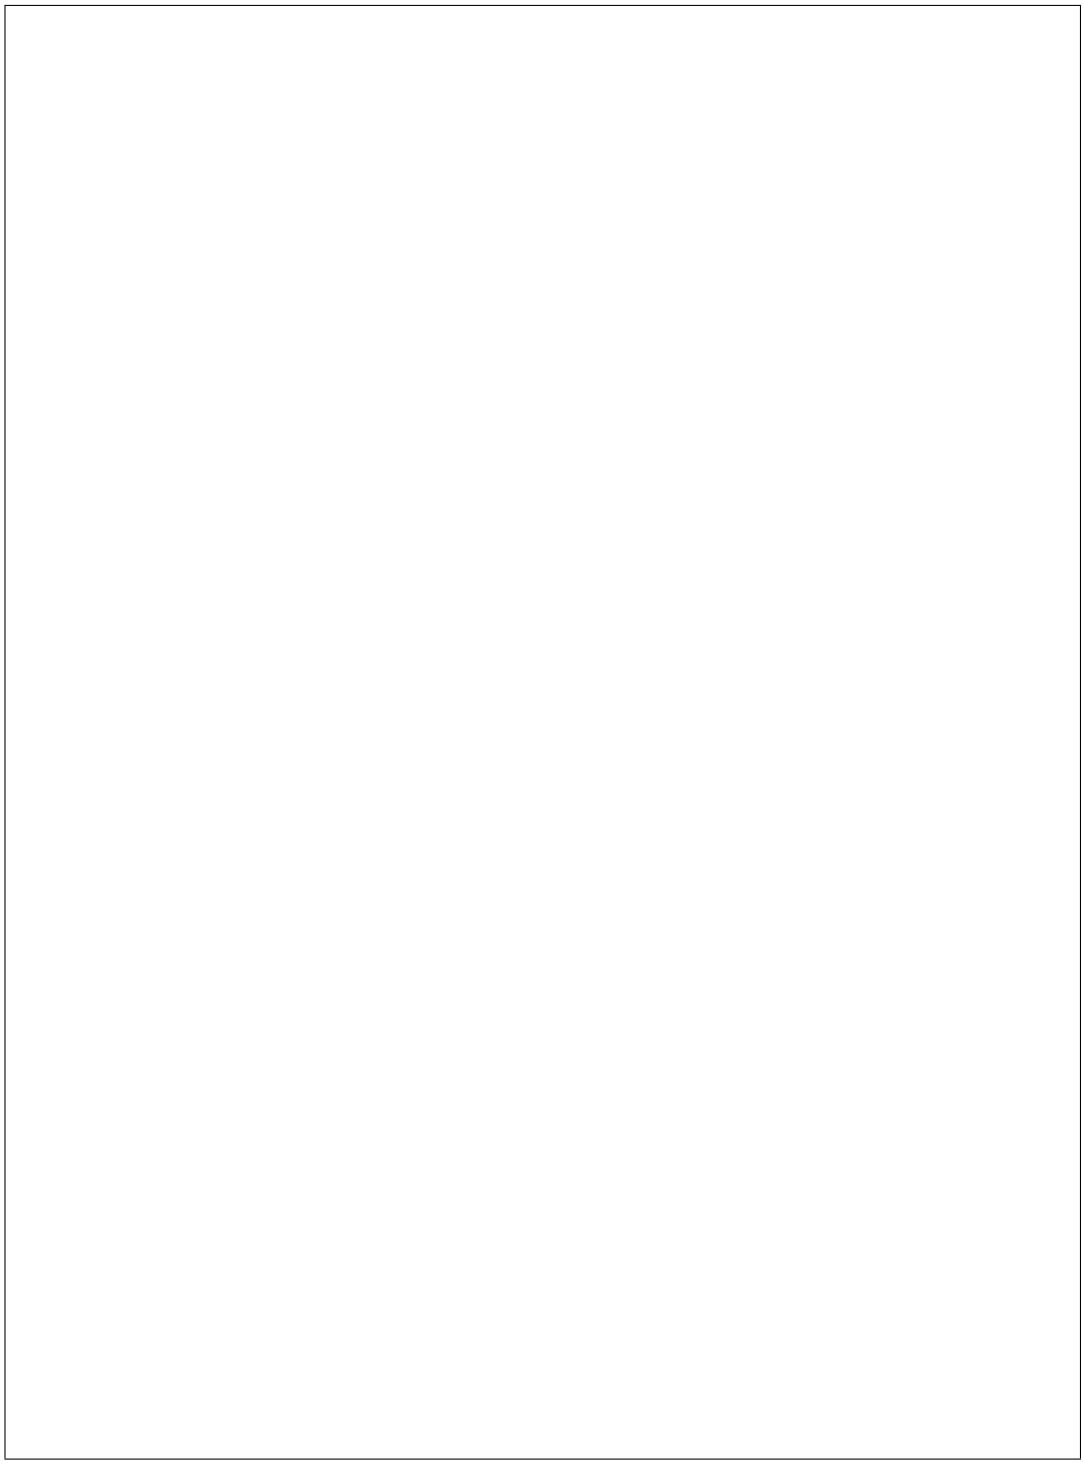

36 Complex matrices

By the end of this section, you should be able to answer the following questions:

- What are unitary, Hermitian and normal matrices?
- Given a complex matrix, determine if it can be unitarily diagonalised, and if so, diagonalise it.

Unitary and Hermitian matrices are complex analogues of orthogonal ($A^{-1} = A^T$) and symmetric ($A = A^T$) real matrices respectively.

In order to define these matrices, we need the following.

36.1 Definition (conjugate transpose)

Let A be a complex matrix. The conjugate transpose of A, denoted A^*, is given by $(\overline{A})^T$, where $\overline{A}$ is the matrix whose entries are complex conjugates of the corresponding entries of A.

Note that if A is real, $A^* = A^T$.

36.1.1 Example

Let $A = \begin{pmatrix} 3+7i & 0 \\ 2i & 4-i \end{pmatrix}$. Write down the conjugate transpose of A.

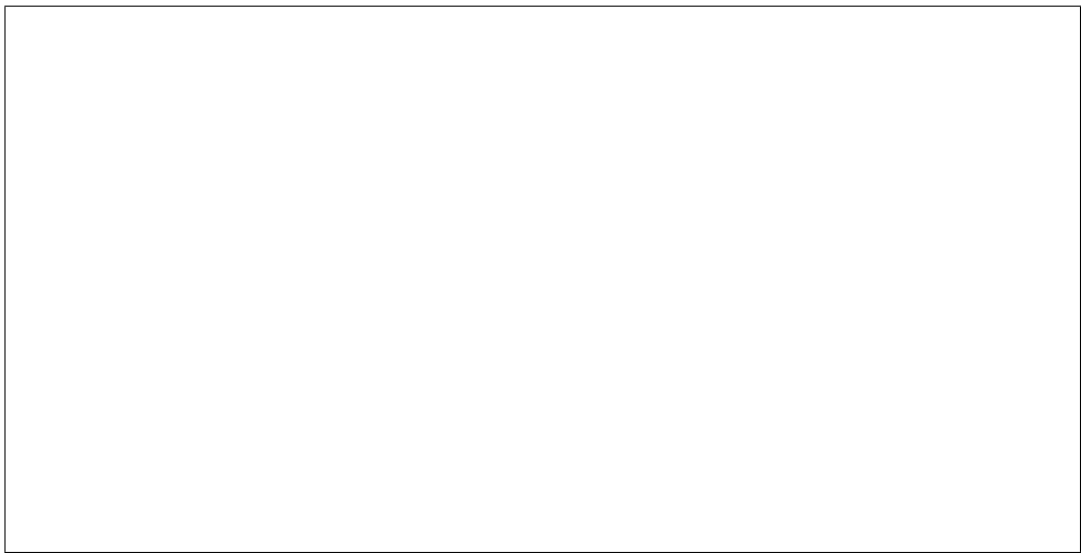

36.2 Unitary matrices

A complex matrix A is said to be *unitary* if $A^{-1} = A^*$. Compare this definition with that of real orthogonal matrices.

Recall that a real matrix is orthogonal if and only if its columns form an orthonormal set of vectors. For complex matrices, this property characterises unitary matrices. In this case however, we must use the complex inner product.

36.3 Complex inner product

Recall that in $\mathbb{R}^n$ the inner (or dot) product of two vectors

$$\boldsymbol{u} = \begin{pmatrix} u_1 \\ u_2 \\ \vdots \\ u_n \end{pmatrix}, \quad \boldsymbol{v} = \begin{pmatrix} v_1 \\ v_2 \\ \vdots \\ v_n \end{pmatrix}$$

is given by

$$\boldsymbol{u} \cdot \boldsymbol{v} = u_1 v_1 + u_2 v_2 + \cdots + u_n v_n$$

and the length (a real number!) of $\boldsymbol{u}$ by

$$|\boldsymbol{u}| = \sqrt{\boldsymbol{u} \cdot \boldsymbol{u}} = \sqrt{u_1^2 + u_2^2 + \cdots + u_n^2}.$$

These definitions are unsuitable for vectors in $\mathbb{C}^n$.

To demonstrate, consider the vector $\boldsymbol{u} = (i, 1)$ in $\mathbb{C}^2$. Using the above expression for length, we would obtain $|\boldsymbol{u}| = \sqrt{i^2 + 1} = 0$, so $\boldsymbol{u}$ would be a non-zero vector with length 0.

Instead, we introduce the complex inner product

$$\boldsymbol{u} \cdot \boldsymbol{v} = u_1 \overline{v}_1 + u_2 \overline{v}_2 + \cdots + u_n \overline{v}_n,$$

where as usual $\overline{v}_i$ denotes the complex conjugate of v_i. In matrix notation, we can write this as $\boldsymbol{u} \cdot \boldsymbol{v} = \boldsymbol{v}^* \boldsymbol{u}$. Note the length of a complex vector is always a real number.

So now we understand what is meant by the following statement: Columns of a unitary matrix form an orthonormal set with respect to the complex inner product.

36.4 Hermitian (self-adjoint) matrices

A complex matrix A is called *Hermitian* (or *self-adjoint*) if $A = A^*$.

As with symmetric matrices, we can recognise a Hermitian matrix by inspection. See if you can see the pattern in the following 2×2, 3×3 and 4×4 Hermitian matrices.

$$\begin{pmatrix} a_{11} & a_{12} + ib_{12} \\ a_{12} - ib_{12} & a_{22} \end{pmatrix}, \quad \begin{pmatrix} a_{11} & a_{12} + ib_{12} & a_{13} + ib_{13} \\ a_{12} - ib_{12} & a_{22} & a_{23} + ib_{23} \\ a_{13} - ib_{13} & a_{23} - ib_{23} & a_{33} \end{pmatrix},$$

$$\begin{pmatrix} a_{11} & a_{12} + ib_{12} & a_{13} + ib_{13} & a_{14} + ib_{14} \\ a_{12} - ib_{12} & a_{22} & a_{23} + ib_{23} & a_{24} + ib_{24} \\ a_{13} - ib_{13} & a_{23} - ib_{23} & a_{33} & a_{34} + ib_{34} \\ a_{14} - ib_{14} & a_{24} - ib_{24} & a_{34} - ib_{34} & a_{44} \end{pmatrix},$$

where $a_{ij}, b_{ij} \in \mathbb{R}$. Note in particular that the diagonal entries are real numbers.

One of the most significant results on Hermitian matrices is that their eigenvalues are real.

36.4.1 Proof that Hermitian matrices have real eigenvalues

Let $\boldsymbol{v} \in \mathbb{C}^n$ be an eigenvector of the Hermitian matrix A, with corresponding eigenvalue λ. In other words,

$$A\boldsymbol{v} = \lambda \boldsymbol{v}. \tag{14}$$

In what follows, we use the fact that $(AB)^* = B^*A^*$ which holds since the same is true for matrix transposition.

We multiply (14) from the left by $\boldsymbol{v}^*$ (treat $\boldsymbol{v}$ as an $n \times 1$ complex matrix) to obtain

$$\boldsymbol{v}^* A\boldsymbol{v} = \boldsymbol{v}^*(\lambda \boldsymbol{v}) = \lambda(\boldsymbol{v}^*\boldsymbol{v}). \tag{15}$$

Also note that

$$(\boldsymbol{v}^* A\boldsymbol{v})^* = \boldsymbol{v}^* A^* (\boldsymbol{v}^*)^* = \boldsymbol{v}^* A\boldsymbol{v}.$$

In other words, $\boldsymbol{v}^* A\boldsymbol{v}$ is also Hermitian. Since it evaluates to be a 1×1 matrix, and all Hermitian matrices have real numbers on their diagonal, this means that $\boldsymbol{v}^* A\boldsymbol{v}$ is a real number.

The quantity $\boldsymbol{v}^*\boldsymbol{v}$ is precisely the complex inner product of $\boldsymbol{v}$ with itself as we have already seen, which is also a real number.

Therefore equation (15) is of the form

$$x = \lambda y, \quad x, y \in \mathbb{R},$$

from which we must conclude that λ is real.

One consequence of this result is that a real symmetric matrix has real eigenvalues, since every real symetric matrix is Hermitian. This result was stated on page 211 but not proved.

36.5 Unitary diagonalisation

We have seen that real symmetric matrices are orthogonally diagonalisable. There is an analagous concept for complex matrices.

A square matrix A with complex entries is said to be unitarily diagonalisable if there is a unitary matrix P such that P^*AP is diagonal.

It is natural to consider which matrices are unitarily diagonalisable. The answer lies in a more general class of matrix.

36.6 Normal matrices

A square complex matrix is called normal if it commutes with its own conjugate transpose, ie, if $AA^* = A^*A$.

Normal matrices are generally more difficult to identify by inspection. However, we have some classes of matrices which are normal:

- unitary,
- Hermitian,
- real skew-symmetric (satisfying $A^T = -A$),
- any diagonal matrix,
- others?

We make a note that real normal 2×2 matrices are either symmetric or of the form $\begin{pmatrix} a & b \\ -b & a \end{pmatrix}$ (which include the skew-symmetric examples).

A class of matrix which is not generally normal is the class of complex symmetric matrices.

36.6.1 Example

Classify the matrix $A = \begin{pmatrix} 1 & 1+i \\ 1+i & -i \end{pmatrix}$.

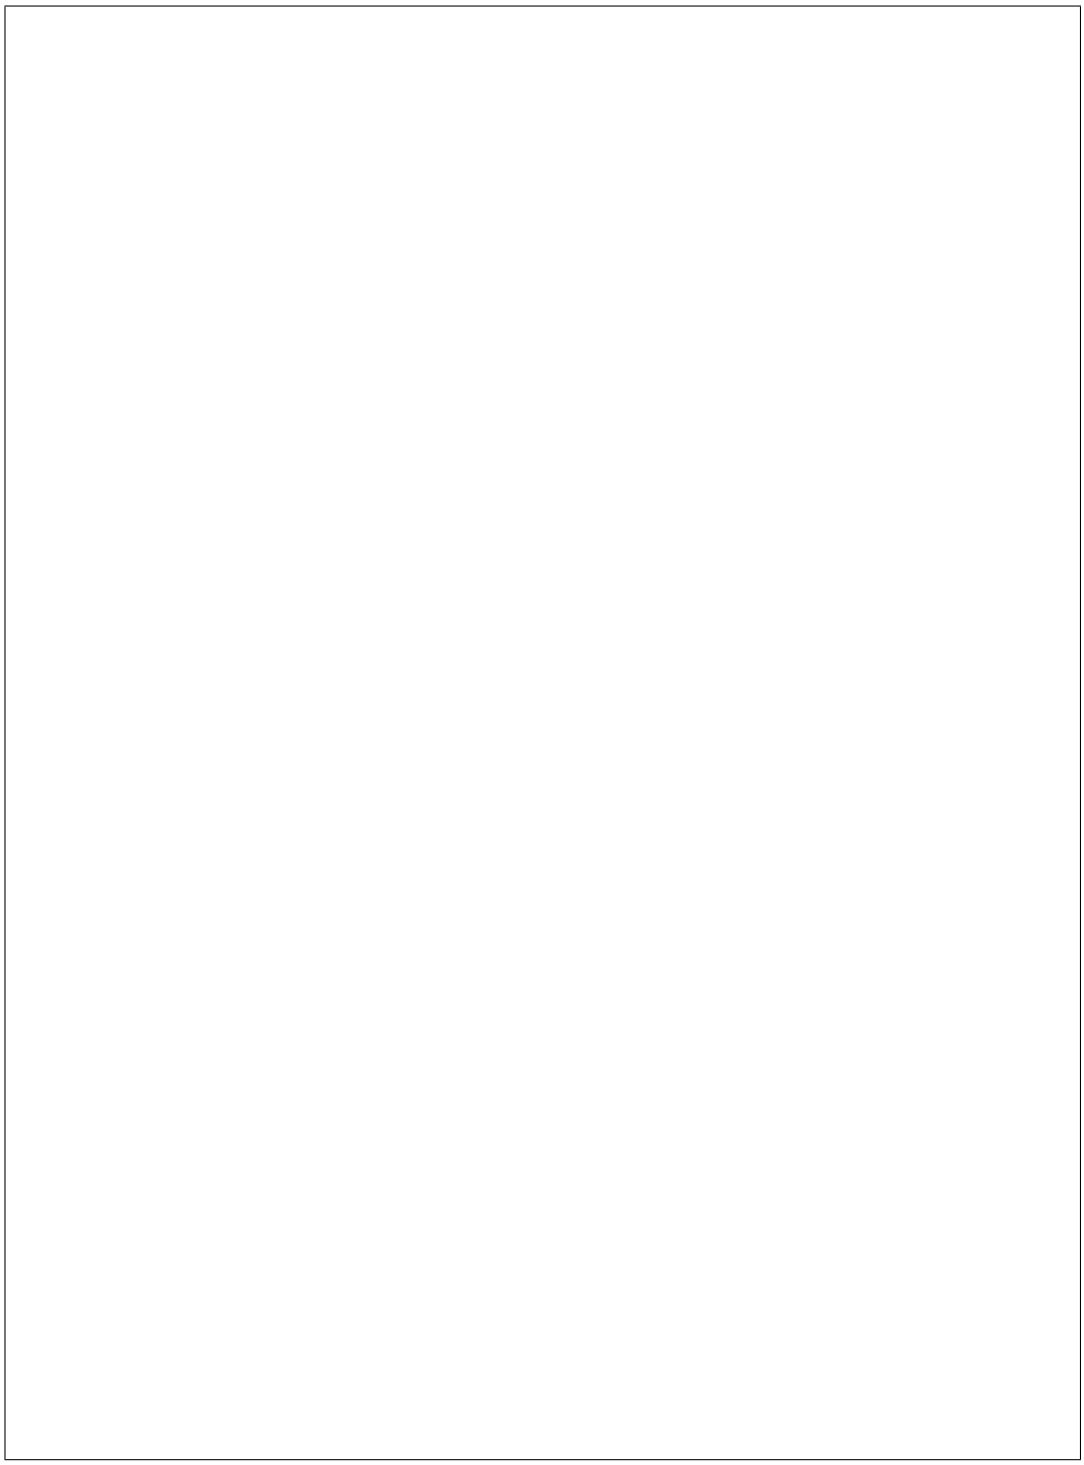

36.7 Normal = unitarily diagonalisable

The main result we have is completely analagous to the real case of orthogonal diagonalisation and symmetric matrices on page 211. We will not prove this result.

An $n \times n$ complex matrix is unitarily diagonalisable if and only if it normal.

36.7.1 Example

If possible, diagonalise the matrix $\begin{pmatrix} 6 & 2+2i \\ 2-2i & 4 \end{pmatrix}$.