1

Introduction

1.1 What this book is about

This book is about applying the concepts of mechanics to the earth and environmental sciences. Mechanics is an old and well-developed branch of physics, which has been applied successfully to many problems in the earth sciences. It is certainly not the only branch of science that has been applied to such problems, but learning how mechanics can be applied to geological problems is probably the best way for a student to learn how to apply other branches of physics to the earth sciences. In the following discussion, we find it useful to make a distinction (not followed by all authorities) between the natural sciences, such as geology or biology, and the physical sciences, such as physics and chemistry.

The natural sciences have made much use of methods developed in mathematics and in the physical sciences. Nevertheless, they have also developed their own unique methods. In geology, for example, students learn to observe in the field and in the laboratory, and to use their observations to give historical explanations of natural phenomena. Mathematics and physics certainly cannot teach us to do this, because they are not much interested in descriptive observation or historical explanation.

Further, all natural sciences are necessarily concerned with particular circumstances, that is, with combinations of materials (e.g., minerals, rocks, plant and animal species) and processes (e.g., the accumulation and flow of ice in continental glaciers, the settling of crystals in magma chambers) actually observed or inferred to be present in the natural environment. Because of the complexity of natural phenomena, the natural sciences have a strong empirical component, and are much concerned with observation and classification. The methodology required to organize this information takes an important and central role in the natural sciences.

The role of the physical sciences, including mechanics, in the earth and environmental sciences can be viewed from two complementary perspectives.

(1) The natural sciences may be seen as an obvious and useful source of applications of basic physical principles. This view fits this book well. The basis of its organization is a (more-or-less) systematic introduction to continuum mechanics. The examples used illustrate the application of these principles to the earth and environmental sciences.

(2) Another view is that mechanics provides a means of explaining many natural physical processes and, more important, a means of generalizing observations from one location to another. Thus, one might attempt to show how the earth's surface can be explained using mechanics. This is not the basis we used to organize the book. Although we have tried to provide examples from a variety of natural environments, we have not attempted to cover all the physical processes acting on, or near, the earth's surface.

Nonetheless, this second view of the role of mechanics in earth and environmental science represents the essential reason why both of the authors, as practitioners in the earth and environmental sciences, have an interest in mechanics and believe that a firm grasp of the subject is necessary for understanding how the earth works and for conducting research that finds utility beyond the time and place of a particular study.

The mechanics of a natural process can be assumed to be universal, so it is not necessary to completely re-explain a physical phenomenon (whether it is river flow, a volcanic eruption, or a landslide) every time it is encountered. This does not mean that to understand one flood, or eruption, or landslide, is to understand them all, because every natural physical process has a setting that is unique. If the essential mechanics is known, however, the basic features of any particular natural event may be more readily understood, unusual aspects may be more easily discovered, and, most important, the nature of the process in the past and future may be assessed.

Much of classical mathematics and physics was developed in order to describe and explain events taking place in three-dimensional space: events such as the motion of projectiles and planets, the bending and breaking of beams and plates, the flow of fluids like air and water, and the slow deformation of crystalline materials like ice. Obviously, these things are interesting not just because they are aspects of how matter behaves, but also because they are a large part of how the earth behaves. To explain any aspect of the earth we need to understand the processes that produced the effects we now see.

Understanding of processes is particularly important in the environmental applications of the earth sciences. In the past geologists were more concerned with understanding history than with making predictions: a geologist spent more time trying to understand how a particular volcano developed over the past few million years, than he did trying to predict how it would behave over the next hundred years. Understanding history is important in the search for minerals and fuels, and it is certainly relevant to the environmental sciences too – but it does not go very far in answering the types of practical questions that the environmental sciences must deal with, such as the following. Will this volcano erupt in the next few years? Will this slope fail? Will the groundwater supply be contaminated by waste disposed at this site?

The physical sciences provide the basic knowledge and techniques needed to answer questions of this type. Generally this type of knowledge has been mostly applied in engineering design: in learning how to construct buildings and bridges that do not fall down, or machinery that operates for some years without serious wear or failure. The environmental sciences are now learning how to combine this type of knowledge about processes with an understanding of historical development in order to make assessments of how the natural environment is likely to change in the future.

This is a challenging and difficult task – and we do not pretend that it can be taught in a single course or learned by reading a single book. But those hoping to practise environmental or earth sciences should at least begin to understand the basic principles, if only so that they can communicate better with experts in the various disciplines involved. Luckily, there is now much that is common to all the major physical sciences, and this core of common knowledge and techniques can readily be applied to the earth. The principles of stress and strain, to which substantial sections of this book are devoted, apply to the deformation of continuous materials, be they fluid or solid, in any environment or application (e.g., the minerals and rocks of the solid earth, natural soils and fluids, the oceans and atmospheres). The laws governing the flow of heat apply to solid rock masses, magma chambers, and lakes, and also have much in common with the laws governing the flow of groundwater. The basic principles of fluid mechanics apply to fluids in all natural environments, and also have close analogues in electrical and magnetic phenomena.

Most of the practical applications of mechanics are based on the assumption that the materials involved constitute a *continuum* in which properties such as mass and strength vary smoothly in space. We discuss this assumption in the next section.

In the remaining sections of this chapter, we discuss four aspects of physical theory that will form themes elaborated in the rest of the book.

(1) In most branches of physical science it is possible to set up *governing equations* which explain and predict the phenomena, at least in principle. Newton's second law relates the change in momentum of a body to the net force acting on it. Though applied by Newton to bodies whose mass could be considered to be concentrated at a point (the "centre of mass"), it took more than a hundred years of research to understand how it could be applied to extended solid or fluid bodies in general. To understand the motion of such materials we need very general formulations not only of Newton's second law (the *equations of motion*) but also of the conservation of mass (the *continuity equations*), and the conservation of angular momentum. We also need to understand how particular materials respond internally to the application of force, a behaviour generally expressed in the *constitutive equations* of that particular material.

(2) Almost all the phenomena of real interest to the earth sciences involve motion in three-dimensional space. To describe such motion, and the forces that produce and resist it, we need appropriate mathematics: it is provided by coordinate geometry, and by vector and tensor analysis.

(3) Once in possession of an appropriate set of governing equations, we need to understand how it can be applied to the problem in hand: this means knowing how to set the problem up by a statement of the *initial and boundary conditions* and how to solve it for those conditions. As the governing equations are generally a set of differential equations (in three dimensions, *partial differential equations*) we need to understand how to integrate these equations. The integrations are generally carried out by a combination of analytical and numerical techniques. Though we cannot explore this aspect of the problem very far, we need to understand the general principles involved.

(4) Besides the science, solving real problems generally involves a good deal of art – something that is generally described as the "art of modeling". This can really only be learned by example and practice, so this book includes many examples of simple applications of mechanics to different branches of the earth and environmental sciences. The examples discussed in the book are necessarily chosen to be as simple as possible, and they may strike some readers as unrealistic. References to more realistic (and therefore more extended and complicated) applications are

given in each chapter. The reader should review at least some of these to get a better sense of the current "state of the art".

1.2 Definition of a continuum

From the point of view of classical physics, matter is generally assumed to take one of three forms: a point mass, a rigid body, or a continuum. All three forms represent an idealization of real matter but the level of idealization decreases as we move from a point mass to a continuum. Though the concept of a point mass is an extreme idealization of real bodies of matter, it is a simplification of the real world that works very well in some applications. It was Newton himself who first showed that we do not need to know the exact distribution of mass within the earth in order to apply the universal law of gravitation. For such purposes, we can assume that all the mass of the earth is concentrated at a point, its centre of mass.

For many other problems concerned with solid bodies, we must be concerned with the size and shape of the material, as well as its mass. For example, the force needed to roll a boulder in flowing water or air depends on the size and shape of the boulder, as well its mass. In this case, we still can simplify the problem somewhat by assuming that the solid body (the boulder) is rigid, even though we know real materials are not perfectly so. Often, the deformation of the solid is very small and if we neglect it we can still obtain interesting and useful results.

The case is not so straightforward when we consider the air or water flowing around the boulder. In doing so, the fluid clearly deforms. That is, there is a change in the position of fluid particles relative to each other as they move around the obstacle. The forces involved in this deformation depend on a variety of factors, including the mass of the fluid and the resistance it offers to deformation. Sometimes we can neglect these properties of the fluid and the associated forces, as we do in Chapter 2, where we first consider the motion of a volcanic bomb without accounting for the drag force imparted to the bomb by the air. Later in that chapter, we add the fluid drag to our discussion and find that the bomb trajectory is considerably different. The difference is caused by forces exerted on the bomb by fluid deformations produced by its movement.

When solids or fluids deform in response to applied forces, we must be concerned with the distribution of material properties (e.g., mass, resistance to deformation) within the material. We still need to make an important simplifying assumption, namely that the distribution of these properties is continuous in space. This is the *continuum hypothesis*. Strictly speaking, this hypothesis cannot be true. We know that the distribution of properties, such as mass, of real materials is decidedly discontinuous at the molecular scale, so we must limit the problems we consider to those with a characteristic length scale several orders of magnitude larger than that scale.

Another problem arises at much larger scales (those typical of many geophysical applications of mechanics) because the properties of a material vary in space. This variation is central to most of the practical problems of concern to us here because it is the gradient, or variation in space, of mass or force that produces motions and deformation. As a result, the useful range in length scale for the continuum hypothesis has two limits: the lower limit is defined by an element of volume that is much bigger than a molecule, and the upper limit is defined by an element of volume that is smaller than any important spatial variation in material properties.

It turns out that such a middle ground exists for most problems. There are of the order of 10^7 molecules in $\approx 10^{-9}$ mm^3 of air. This is both a huge number of molecules (so we need not worry about the exact number of molecules in an elemental volume) and a very small elemental volume (much smaller than the physical dimensions of any practical problem). A limiting volume of 10^{-9} mm^3 is applicable to all liquids and to gases at atmospheric pressure.

Further, consider the density of a fluid. In the natural sciences, we rarely consider discrete "bodies" of fluid, so that it is usually difficult to define what is meant by the total fluid mass in a particular case. It is easier to measure the mass contained within a given volume of the fluid. Then we can define the mass density (or simply, density) ρ, as the measured mass divided by the measured volume. Mass is a property of points (or bodies considered as point masses); density is the corresponding property of a continuum. To make the concept more precise we reduce the measured volume to "elemental" size. As long as our elemental volume of fluid is much larger than the lower scale limit, so that it contains a huge number of molecules, we need not be concerned with exactly how many molecules, or even which molecules, are in the volume and we can define the density with reasonable accuracy. As long as the elemental volume is also much smaller than the upper scale limit over which fluid density varies appreciably, the concept of density can be usefully applied to the real world.

The minimum length scale for continuum analyses is not necessarily the molecular scale. For example, consider the density and strength of a pile of cobbles the size of baseballs. For a test volume of the order of a cubic millimetre, the density would vary wildly between that of air and that of the rock cobbles, depending on where the volume was located. On the other hand, the density within a volume of many cubic metres would show only negligible variation with the exact placement of the test volume. Similarly, if we wished to measure the aggregate strength of the cobbles, it is reasonable that the scale of the testing apparatus would have to be at least two orders of magnitude larger than an individual cobble to avoid the variation that would result from placing different numbers and arrangements of cobbles in the test volume. These issues of a minimum length scale apply not only to solids such as soils and rock, but also to mixtures of fluids of different composition (e.g., oil and water), and mixtures of fluids and solid particles (e.g., slurries or suspensions of sediment in water).

The reason we use the continuum hypothesis is simply that it allows us to use differential calculus to analyze the properties of a material and its motion and deformation. Continuous changes, or *gradients,* in physical properties and forces define mechanical problems, and differential calculus is the mathematics that treats such gradients precisely and efficiently. In using the continuum hypothesis, we are not limited to cases without abrupt boundaries between different materials. In these cases, we have to define the appropriate *boundary conditions* and then we can use continuum mechanics to describe what happens within those boundaries.

1.3 Governing equations

In Chapter 2, we shall briefly review the mechanics of points and extended bodies. We show how Newton's second law can be regarded as the governing equation for the motion of point masses. It relates the observed change in momentum $d(mv)/dt$ to the sum of all the forces $\sum F_i$, acting on a point mass, m:

$$\frac{d(mv)}{dt} = \sum F_i \tag{1.1}$$

Equations like this, generally written out for each of the three coordinate directions separately, are called the *equations of motion.* They are found not only in the classical mechanics of point masses, but also in the mechanics of rigid bodies, and in the mechanics of continua. In the mechanics of rigid bodies, applied forces not only produce accelerating motion in the direction of the mean force ("translation"), but may also produce accelerating rotations. In the mechanics of continua, applied forces may also produce distortion, that is, elongation or contraction, and shear.

For the motion of point masses, the governing equations consist of Newton's three laws, together with the (generally implicit) assumption of conservation of mass – that is, the mass does not change with time. But for the dynamics of rigid bodies we need to add further concepts, those of torque, or the moment of a force, and angular momentum (the "moment of momentum") which were not stated by Newton, and indeed were not explicitly formulated until almost 80 years later by Euler (see Chapter 2).

For continua, the situation is much more complex. To start with, we have to consider force in a different way. In studying the application of a force to a point mass, we are not faced with the problem of how the force varies over the surface of the mass. But in studying the deformation of continua, it is not only the total applied force that matters but also how it is distributed over the surface to which it is applied. We are concerned, therefore, with forces per unit area, or stress, and how the stress changes across a small volume of the medium. Thus, it is the gradient of the stress that matters, and so the equations of motion are generally written in units of force per unit volume, rather than force as in Equation (1.1). (The derivative of a stress, that is a force per unit area, with respect to distance has the units of force per volume.)

Even after converting from forces to stress gradients, the governing equations for a continuum involve more than the equations of motion and continuity. These equations tell us how the stresses act on the continuum, but they tell us nothing about how a deformable continuum responds to the applied stresses. Consider small cubes of rubber, clay, and motor oil. These different materials will respond quite differently to the same applied stresses. Indeed, one fundamental difference is immediately evident: we need a container to keep oil in the shape of a cube whereas the rubber and clay are able to stand freely. The equations of motion do not explain the various responses of different deformable materials to applied stresses; another set of equations, each specific to a particular type of material, is needed. These are the *constitutive equations* for that particular continuum.

Fortunately, it turns out that three simple constitutive equations do a good job of describing the response of many natural materials to applied stresses. Rubber, clay, and motor oil are examples of materials represented by each type of equation. Each constitutive equation is also conventionally represented by a simple mechanical model.

(1) Rubber is an example of the first type of material: it is a *linearly elastic* solid. When a stress is applied, its deformation is finite and directly proportional to the applied stress. The deformation is essentially instantaneous and is completely recoverable when the stress is removed. The constitutive equation for an ideal elastic substance is a linear relation between stress and strain. After the mathematics of material deformation

(strain) are presented in Chapter 7, the behavior of linearly elastic solids is discussed in Chapter 8. The conventional mechanical model is a spring that is fixed at one end. When a force is applied to the other end, the elongation of the spring is directly proportional to the force. When the force is removed, the spring recoils to its original position.

(2) Clay represents a *plastic* material. If the applied stress is too small, the clay does not deform at all. Once the level of applied stress reaches some threshold strength, however, the clay deforms permanently. For an ideal plastic, the applied stress cannot be increased beyond this threshold value, because it is immediately relieved by flow. The constitutive equation for an ideal plastic is very simple and states that the stress within the body is less than or equal to the threshold strength. If a stress larger than this is applied, it is immediately reduced to the threshold stress level by flow. Plastic behavior and the idea of material strength are covered in Chapter 4, once the essentials of stress have been presented. The conventional mechanical model is a block resting on a flat, rough surface. The block does not move until the force applied is sufficient to overcome the friction between the block and the surface below. The deformation in a plastic material is permanent, unlike that in an elastic solid. When the force is removed, the block does not return to its original position.

(3) Motor oil represents a *viscous* substance. Viscous materials strain indefinitely in response to an applied stress and, for an ideal viscous fluid, the constitutive equation is a linear relation between the applied stress and the rate of strain. The behavior of linear, or Newtonian, viscous fluids is covered in Chapter 9. The conventional mechanical model is a *dashpot*, which is a small cylinder filled with oil. The cylinder contains a loose-fitting piston that can be pushed back and forth within the cylinder. As the piston is moved, the fluid moves from one side of the piston to the other at a rate that is proportional to the force applied to the piston. In contrast to an elastic substance, but similar to a plastic material, a fluid deforms permanently: when the force is removed from the piston, the fluid and piston remain in place. Unlike a plastic substance, but similarly to an ideal elastic material, a viscous fluid will deform at any stress, regardless of its magnitude.

Ideal elastic and viscous materials have linear relations between stress and strain or strain rate. The stress applied to an ideal plastic can be no larger than the threshold strength. Although the behavior of many natural materials can be approximated using these three simple models, many others show more complex behavior. Still, relatively simple modifications to the basic constitutive equations can be made to approximate the stress–strain behavior of most natural materials. These include nonlinear elastic or viscous constitutive equations and simple, additive combinations of the basic relations for elastic, plastic, and viscous behavior. We will consider some of these in Chapter 10.

1.4 Vectors and tensors

In elementary mechanics, a familiar concept is that of a *vector*, which can represent a quantity like force, velocity or acceleration, that has both magnitude and direction. But a force per unit area, called a *stress*, can only be defined by specifying a magnitude and *two directions*, the orientation of the force itself, and of the surface on which it acts. As continua transmit stresses to all parts of the body, a stress applied to one surface can be felt as a stress on any other surface within the body. So if we want to express the complete *state of stress* at a point we must generally specify the stresses, not just on a single surface but on three surfaces normal to the three coordinate axes. We will see in Chapter 4, that stress must therefore be described by an array of six components, not just the three components that are sufficient for force.

The simplest stress field can be described by a single scalar quantity, pressure. Some of the applications of pressure are, however, far from simple, and in Chapter 4 we introduce the important concepts of buoyancy and effective stress. The consideration of pressure in pore spaces leads naturally to the consideration of flow through porous media, in Chapter 6. An understanding of pore pressures and the movement of pore fluids is important, not only for practical applications in soils, geohydrology, and petroleum geology, but also in understanding the stress, strength and deformation (e.g., compaction) of porous soils and rocks.

Stress is by no means the only quantity that is more complicated than a vector. When we study how materials deform (Chapter 7), we will see that to describe deformation we need a quantity, called *strain*, which is at least as complicated as stress. And the physical properties of solid continua, such as the way they transmit fluids (their *permeability*, Chapter 6) or respond to stresses (their elastic or viscous moduli) may also be complicated quantities. Just as a special notation and algebra, *vector algebra*, was developed to describe vectors, so a special notation and algebra, called *tensor* algebra, has been developed to describe tensor quantities such as stress or strain. In this book we do not attempt to provide a complete discussion of either vector or tensor algebra, but we do introduce the notation (Chapters 2 and 4), and describe a few of its uses. Vector operators (grad, div, curl), essential for the description of deformation (Chapters 7 and 8) and flow (Chapter 9), are introduced in Chapter 6, in the simpler context of flow through porous media.

One of the great virtues of tensor notation is that by using it, we can write complicated results very compactly – so it is widely used in the professional literature, and without some understanding of the notation it is impossible to read the journals that describe the applications of mechanics to problems in the earth sciences. We also come to see that vectors can be regarded as a special kind of tensor: a "first-order" tensor, that needs only a single row or column of numbers, unlike the "second-order" tensors generally met in mechanics, which require square arrays of numbers for their description. Ordinary scalar quantities, like mass, that require only a single number, can in turn be regarded as "zeroth-order" tensors (Chapter 4).

The concept of stress, though it now seems to us a rather obvious extension of the idea of a force, was the product of almost a century and a half of thought by the best mathematicians and physicists of the time, and was first clearly formulated by Cauchy in 1822. It is not surprising, therefore, that most of the progress in formulating the governing equations of continuum mechanics took place in the nineteenth century: Newton himself had only a very hazy idea of what constituted a Newtonian fluid.

1.5 Solving the equations

Though Newton formulated the basic equations of motion, and derived many interesting results from them, he did not introduce the formal techniques now used to apply the governing equations to particular phenomena. This was done over the following 50–100 years by a group of Swiss and French scientists and mathematicians. The pattern now in almost universal use is to express the governing equations as sets of partial differential equations that state how the equations of motion, the continuity equation, and the constitutive equations apply to each infinitesimal "element" of matter. The problem is then to integrate these equations (generally more than once) taking into account the particular configuration of interest. So, for example, if we want to know how water flows through a channel, we need to know the *initial conditions*, that is, the depth and velocity of the water at the first time that interests us, and the *boundary conditions*, for example, the geometry of the channel, and the specifications that the velocity must fall to zero at the solid boundaries of the channel and that the shear stress is zero at the free surface. Only then can we hope to discover, for example, how the stresses vary along the channel floor, or how the velocity is distributed within the flow.

Integrating the governing equations is generally possible only for a few simple initial or boundary conditions. This does not mean the equations are useless: even if they cannot be integrated, study of the equations often reveals many interesting general properties of the phenomena – and in particular, it tells us how they can be applied at different geometric and temporal scales (Chapter 3). In fact, the scaling properties of general physical laws, first discussed thoroughly only in the mid-nineteenth century, has proved to be a very powerful general method applied more recently to many complex problems in nuclear and solid state physics, as well as to problems such as the turbulent flow of fluids (Chapters 11 and 12).

Since the development of digital computers, enormous progress has been made in the numerical solution of differential equations. The subject is a large one, but simple examples are certainly not beyond the capacity of the personal computers now readily available to students. We have prepared a set of computer programs that supplement the material presented in this book, and illustrate a few of the techniques that are now routinely used in almost all applications of mechanics to real scientific and engineering problems.

1.6 The art of modeling

In Chapter 3, we discuss at some length just what is meant by a model in a modern science such as mechanics. Models may be actual physical models of phenomenon, generally built at a reduced scale, or they may be conceptual or mathematical constructs. Constructing models generally means greatly simplifying the real world, but it also means making assumptions that bridge over the gaps in our knowledge. It is useful as an intellectual exercise, since it forces us to confront both our ignorance and also how we can apply our limited knowledge to a real situation; and it is useful for making predictions.

Almost all environmental questions demand the construction of models. Will burning fossil fuels cause global warming? We need a model of the fate of carbon dioxide released into the atmosphere by burning fuel, and another model of how this will affect the global climate. Yet further models are needed to predict how climate change would affect sea level, or energy consumption, or the growth of crops. It is impossible to predict exactly what will happen – the natural systems involved are much too complicated for any model to give a complete description of what is likely to happen. A model attempts to include what is known about all the major factors and to assess their relative importance. It is indeed helpful to discover and make use of the historical background – we could make little progress in understanding the present climate if we did not know from geological history that we have only recently emerged from a major ice age – but history is not enough. We must also try to understand the physical processes that govern natural phenomena, and how (or even if) these can be reliably applied to predict what will happen next.