1

INTRODUCTION

Molecular dynamics simulations compute the motions of individual molecules in models of solids, liquids, and gases. The key idea here is *motion*, which describes how positions, velocities, and orientations change with time. In effect, molecular dynamics constitutes a motion picture that follows molecules as they dart to and fro, twisting, turning, colliding with one another, and, perhaps, colliding with their container.

This usage is not unique: *molecular dynamics* may also refer to the motions of real molecules when studied primarily by molecular beam [1] or spectroscopic [2] techniques. This terminological confusion is compounded by *lattice dynamics* [3], which refers to the study of vibratory motions of atoms in solids, and by *molecular mechanics* [4], also called force field calculations, which refers to quantum mechanical calculations of the structure of individual molecules. This book is concerned with molecular dynamics solely in the sense of simulation.

Molecular dynamics simulation is the modern realization of an old, essentially old-fashioned, idea in science; namely, the behavior of a system can be computed if we have, for the system's parts, a set of initial conditions plus forces of interaction. From the time of Newton to the present day, this deterministic mechanical interpretation of Nature has dominated science [5]. In 1814, roughly a century after Newton, Laplace wrote [6]:

> Given for one instant an intelligence which could comprehend all the forces by which nature is animated and the respective situation of the beings who compose it—an intelligence sufficiently vast to submit these data to analysis—it would embrace in the same formula the movements of the greatest bodies of

> the universe and those of the lightest atoms; for it, nothing would be uncertain and the future, as the past, would be present to its eyes.

If this approach is thwarted by the complexities of reality, then we replace reality with a model. In one of his Baltimore lectures (Lecture XI), roughly a century after Laplace, Thomson observed [7]:

> It seems to me that the test of "Do we or not understand a particular subject in physics?" is, "Can we make a mechanical model of it?"

Today, roughly a century after Thomson, we remain undeterred from Laplace's dream: the requisite "intelligence" is provided by the digital computer, the "respective situation" is a set of initial positions and velocities, "the same formula" though not literally true could be interpreted as the same algorithmic program, and Laplace's universe has given way to model universes. Now, deterministic mathematical models pervade not only the physical sciences and engineering, but the life and social sciences [8] as well.†

This attitude is old-fashioned in the sense that, while often successful, it is nevertheless simplistic. In spite of Laplace's claim, we can still identify systems that are unpredictable—stock markets and the weather, for example. Why should this be? If deterministic mathematical models can help us successfully land *Apollo XI* on the moon, why can't they help us predict next month's weather on earth?

The resolution of this dilemma is based on the kind of forces acting among system components: when a system contains objects that interact *nonlinearly*, the system's behavior may be unpredictable. In the past few years studies in nonlinear dynamics have decoupled deterministic from predictable [9]. *Deterministic* situations have system outputs causally connected to system inputs. *Calculable* situations are those deterministic situations in which an algorithm allows us to compute system outputs if the inputs were known. *Predictable* situations are those calculable situations in which the algorithm can be numerically implemented to actually compute the outputs. Calculable situations may be unpredictable because of the large number of inputs needed, because of an unrealistically high precision with which the inputs must be known, and/or because the algorithm's stability is sensitive to intermediate calculations. In pool, *Eight ball in the side pocket* is deterministic, calculable, and predictable; however, whether it will rain in two weeks is deterministic but unpredictable.

The overriding theme of this book is predicated on the decoupling of predictability from determinism. Be warned—that you use a machine to

†In fairness, paleontologists, at least, have discovered deterministic unpredictability. Thus, Stephen Jay Gould [10] posits that if the tape of life were rewound to some previous, sufficiently removed condition and then replayed, the result would be life unlike life as we know it. For more technical conjectures on connections between life and deterministic unpredictability, see Fox [11].

compute the behavior of a many-body model does not guarantee that the computed behavior is representative of that model, much less that the model mimics reality. To my mind it is this deterministic unpredictability that makes molecular dynamics fascinating and challenging. Is the fun (aka intellectual stimulation) merely in making a model, writing some differential equations, and loading it all into a computer? No. The fun, it seems to me, begins when we have completed a simulation, when we have a number. Now there arise all the old familiar questions characteristic of science: How good is this number? How could it be wrong? Is it representative, that is, reproducible? What does it mean? Do I believe it? How do I test it? If it is right, what must follow? This book should not only help you learn how to simulate but also make you aware that questioning the results is part of the procedure.

Computer simulations are performed on models, not on real things, and so the science of simulation, while distinct from, is necessarily bound to the art of model building. The purposes of this chapter are to clarify the distinctions between models and simulations and to discuss how together they contribute to new understanding.

1.1 SYSTEMS AND ALL *THAT*

The portion of the physical world on which we focus our attention is called the *system*; it is a subset of the universe. The system may be composed of any number of similar or dissimilar parts and the condition of those parts identifies the *state* of the system. For example, the door to my office constitutes a system to which I can ascribe two states: *open* or *shut*. To analyze and describe the behavior of the system, we need ways for assigning numerical values either to the state or to functions of the state; such assignments are called *observables*. Thus, to my door we can ascribe an observable called openness, to which I assign the value 1 if the door is open and 0 if the door is shut. As another example, let the system be 10^{24} molecules of a gas. Its state is specified by the position and momentum of each molecule, and the state gives rise to such observables as temperature and pressure.

The state of a system can be manipulated and controlled from the environment via *interactions*. For example, I may change the state of my office door by an interaction, specifically, by exerting a force. Allowed interactions are constrained by the nature of the boundary that separates the system from its surroundings. Various kinds of boundaries and interactions are possible, but in this book we will limit our attention to *isolated systems*: systems that can exchange neither matter nor energy with their surroundings.

We cannot usually study a system by directly observing the state; instead, we probe states indirectly by manipulating, controlling, and measuring observables. Thus, studies of a gas may involve controlling the system volume, manipulating its temperature, and measuring how the pressure responds. The

isolated system is special in the sense that we do not interact with it, and therefore we can manipulate its observables only before the system is isolated.

To organize, describe, and perhaps even predict observables, we create theories. Theories may operate at one of several levels. At the simplest level, theories merely provide relations among observables. For example, the ideal-gas law

$$PV = NkT \tag{1.1}$$

was originally obtained by organizing results from measurements of the pressure, volume, and temperature of low-density gases.

At the next level of complexity are theories that relate observables to the underlying state. For example, at this level we have kinetic theory, which teaches that the observable temperature is related to the state through the molecular velocities. Theories at this level provide interpretations or explanations for observables, but if the state itself is unobservable, these theories cannot be used to compute numerical values for observables.

To overcome this computational dead end, two strategies have been devised: (1) concoct theories at still higher levels or (2) perform computer simulations. Higher level theories try to resolve the computational difficulty by reorganizing and reducing the detailed information about the state needed to compute values for observables. Such is the objective of statistical mechanics, in which observables are related not to the underlying state itself, but rather to the probability of the system being in particular states.

The alternative strategy includes molecular dynamics. Molecular dynamics assigns numerical values to states, thereby making states observable, at least for model substances. With numerical values assigned to states, theoretical relations from kinetic theory can be used to compute values for experimentally accessible observables. Thus molecular dynamics is closely tied to kinetic theory and not as closely related to statistical mechanics.[†] In particular, molecular dynamics is less sophisticated, less elegant, but more direct than statistical mechanics.

1.2 MODELING VERSUS SIMULATION

Whether we study systems theoretically or experimentally, the general procedure is the same: we manipulate and control certain observables (inputs), the

[†]W. Thomson, Lecture I, p. 1: "...the kinetic theory of gases is a part of molecular dynamics, is founded upon molecular dynamics, works wholly within molecular dynamics, to it molecular dynamics is everything, and it must be advanced by molecular dynamics..."

system responds, and then we measure or compute other observables (outputs):

Since theoretical analyses are now largely done via modeling or simulation or both, it is instructive to clarify how modeling differs from simulation. For a discussion more complete than what follows, see Casti [8].

The goal of theoretical work is to establish connections between measurable outputs and controlled inputs. In Section 1.1 we discussed how this may be done at different levels of complexity; in particular, sophisticated theories use an underlying state to connect outputs to inputs. Part of the theoretical problem is to define the state in such a way that complicated interactions among state variables are decoupled, or at least weakened, so that observable outputs can in fact be computed.

A model is an attempt to decouple and remove interactions that have little or no influence on the observables being studied. Thus, a model is simpler than the system it mimics: it has access to fewer states. Decoupling interactions means relaxing constraints; hence, a model has access to some states not available to the original system and vice versa. In other words, a model is a subset or subsystem of the original system: outputs from a model will be consistent with those of the original system, but only for a restricted set of inputs. For those restricted inputs, since the model is a subsystem of the original, states visited by the model correspond to those visited by the original system.

In contrast, a simulation is more complicated than the system it simulates: a simulation generally can reach many more states than can the original system. A simulation imposes constraints so that the simulated output is consistent with the output of the original system, at least for a restricted set of inputs. A simulation will typically bear no structural relation to the original system; for example, the way constraints are imposed in the simulation may differ from the mechanism that confines the original system to certain states. Hence, states in the simulation may bear no correspondence to states of the original system. Although a simulation is more complex than the original system, it does *not* follow that the original system is a model of the simulation.

An example should clarify these ideas. As the real system, consider a simple ball-and-spring arrangement. One end of the spring is attached to the ball, the other end is fixed to a wall. In response to a displacement from its equilibrium position, the ball slides on a floor. The problems are to, in turn, model and simulate the motion of the ball that results from a displacement.

Let R represent this real system, that is, the spring as mover plus the ball as the thing moved.

To study the motion of R, we might construct a device M that is a one-dimensional harmonic oscillator (ODHO) having spring constant γ. The motion of M is described by the differential equation

$$\frac{d^2x}{dt^2} = -\gamma x(t) \tag{1.2}$$

where x is the distance of the ball from its equilibrium position and t is time. Writing this equation presumes several simplifying assumptions: (a) the motion of the ball is restricted to a line, (b) the spring is perfectly harmonic, (c) the ball experiences no sliding friction on the floor, and (d) the ball has no internal states that exchange energy with the spring. These assumptions imply that M is a simple subsystem of R. For some initial displacements M will mimic the motion in R; however, M cannot mimic all the behavior available to R. For example, in R we might initially raise the ball from the floor, allowing the ball to move in the xz-plane, motion not allowed to M. The one-dimensional harmonic oscillator M is a *model* of the real system R.

An alternative scheme S for studying the motion of R would be to remove the spring and use a person as the mover. This person might be a well-trained graduate student who has the uncanny ability to move the ball, for many sets of initial conditions, so as to reproduce the motion of the ball in R. This situation S is more complex than R because a person is more sophisticated than a spring. Moreover, S involves imposing constraints on the student's arm to make the motion of the ball mimic the motion in R. Otherwise, more states are possible in S than in R; for example, the student might absent-mindedly drop the ball in his pocket when he stops for coffee. The situation S is a *simulation* of the real system R.

Note that as well as simulating the real system R, we could also simulate the model M. For example, we might have another mover (a professor or, equivalently, a robot) that can move the ball through states visited by the perfect ODHO model M. This simulation of M is more complex than M itself, but the added complexity does not make the model more realistic, that is, more like the real system R.

How do these distinctions relate to molecular simulations? Well, what do we typically do? We identify a substance and its observables that we want to study, say, thermodynamic properties of argon. Then we construct a model of the substance, say, the spherically symmetric, pairwise additive Lennard-Jones potential. This is a true model. The Lennard-Jones potential is simpler than the argon potential because argon atoms are not perfect spheres and their interactions are certainly not only pairwise additive. With the model chosen,

we then perform a simulation—but a simulation of what? It can only be a simulation of the model, of the Lennard-Jones substance. We do not simulate argon. The simulation is more complex than the model, but the added complexity does not add to the realism of the resulting observable outputs. In error are those who claim that molecular dynamics simulates argon, or water, or proteins, or whatever. We simulate molecular models of such substances.

1.3 THEORY VERSUS EXPERIMENT

At a scientific conference in the 1970s there broke out a heated debate as to whether computer simulations like molecular dynamics are theories or are experiments. The theory side argued that simulation is clearly not experiment because no measurements are done on real systems; molecular simulations are pure calculation. The experiment side countered that simulation results are *used* like experiments, namely, to test theories; it isn't sensible to test one theory with another theory is it? Moreover, this side noted that simulation results, like experiments, are prone to problems of reproducibility and statistical error. Hence, pervading the literature is the interpretation of molecular simulations as computer experiments.

What's the resolution of this dilemma? And does it really matter how we think of simulation? Consider the following example. To perform an experiment, to take a measurement, the observer must interact with the system: some type of probe necessarily has to cross the system boundary. Thus, truly isolated systems cannot be studied experimentally: once our probe crosses the boundary, the system is no longer isolated. However, we can perform theoretical calculations, such as simulations, on truly isolated systems and obtain meaningful results. The resolution of our dilemma has to be that molecular simulations are forms of theory; they do not involve measurements on real systems.

How we think of simulation is important, indeed crucial, because of the consequences: if we accept that simulations are experiments, then it follows that the models simulated are real. Armed with this attitude plus the ease of actually doing simulations, we may be tempted to abandon laboratory experiment altogether. The danger lies in severing simulation from reality.

Is this only an academic issue? Recently a public presentation was made by the chief executive officer of a major computer manufacturer. In the talk *Voyager* photos of Jupiter were compared with images produced from a computer simulation of Jovian weather. The speaker's punch line was that the simulated images were actually "more real than real life." But if this were true, then there would be no need for further explorations by spacecraft —we would simply perform simulations. Note that, in fact, reality doesn't enter this picture at all. The CEO can be faulted on two counts: not only did

he erroneously claim that a simulation can supplant reality, but he also confused a photographic image with reality.†

To illustrate this point in a more mundane situation, consider: Police have been summoned to the scene of a domestic quarrel. In the kitchen the patrolmen find a bewildered wife standing over the prostrate form of her husband.

"Okay, lady, what happened?"

"I don't rightly know. Herb riled me, so I hit'em in the head with a tomato."

"Sure, lady. There's no tomato on the body or the floor."

"Isn't there? O'course not. T'was a decorative ceramic."

And lest you feel that this contrived example‡ is only impractical philosophical quibbling, consider the intense ethical debates that were prompted by American television broadcasts of selected news events in the form of interpolative reenactments (aka *simulations*).

1.4 REDUCTIONISM VERSUS SIMULATION

Since the time of Newton, scientific theories have nearly all been developed in a reductionistic mode: a complex system is reduced to one or more simple subsytems and the subsystems are analyzed. Subsystems ultimately take the form of models, and today models are almost exclusively mathematical. Before about 1960 mathematical models had to be simple enough to be tractable analytically, but now this constraint is relaxed by the availability of digital computers.

As an example of reductionism, consider the study of matter in simpler and simpler forms:

$$\text{Matter} \rightarrow \text{compounds} \rightarrow \text{elements} \rightarrow \text{molecules} \rightarrow \text{atoms}$$

$$\rightarrow \text{elementary particles} \rightarrow \text{quarks}$$

Modeling may occur at any stage of such a reduction. Successful modeling requires a construction that forces the behavior of interest to remain *invariant* when the subsystem is replaced by the model. The goal of reductionism is

†A principal feature of science is the apparently endless disentangling of images from reality. Well over 800 years ago in his *Questiones Naturales*, the scholar Adelard of Bath was moved to write, "Wherefore, if you want to hear anything more from me, give and take reason. For I'm not the sort of man that can be fed on a picture of a beefsteak" [12].

‡This example is a slight modification of a vaudeville routine used by James Thurber [13] to illustrate *confusing the container with the thing contained.*

to explain system behavior by combining explanations for the behaviors of its subsystem models. Simulation provides an alternative to reductionism because simulation allows us to study the behaviors of classes of systems or subsystems. Thus, while reductionism emphasizes structural analysis, simulation emphasizes behavioral classification. As an example, consider study of the fluid–solid phase transition. One simulation approach to this problem would be to load marbles into a drum. To force the marbles to move, we rotate the drum (imagine a cement mixer). As the drum rotates, we add more and more marbles, until finally enough are added to freeze the motion. From the number of marbles N and the drum volume V, we obtain the density N/V needed for solidification. We then repeat the simulation with balls of other diameters: golf balls, baseballs, soccer balls, basketballs. By inspecting the solidification densities for balls of various diameters, we conclude that one mode of solidification is a geometric packing effect, controlled by the packing fraction $V_{\text{balls}}/V_{\text{drum}}$. Note the features of this simulation: the controlled input observables are the ball diameter, the number of balls, and the drum volume; the measured output observable is the density at which motion ceases. During the experiments, the states of the balls—positions and velocities—bear no relation to states of molecules in any real substance. Further, the entire study is not just of one substance, but rather a systematic progression through a class of substances: spheres of increasing diameters.

In contrast, a reductionist would study this problem by combining a model with a theory to predict the solidification of a particular substance. The theory might involve only relations among observables, such as *PVT* equations of state, or the theory might include underlying system states, such as is done in statistical mechanics. In any case, the validity of the model and the theory would be tested by comparing predictions with experimental measurements on a real substance. However, the connection between input and output observables would remain *implicit* in the mathematical apparatus used to make the prediction.

But the goal of the simulations is not so much to predict solidification as to make *explicit* how input and output observables are connected. In other words, rather than predictions, the goal is more in the nature of providing explanations: idealized models "*explain* nature even while they do not describe it" [14]. Both reductionism and simulation contribute to science; however, for a specific problem one or the other may be more appropriate. In particular, simulation is not always the best method.

As shown in Figure 1.1, we identify two distinct roles that simulation can play in scientific investigation. At the higher level, simulation, including computer simulation, serves as an alternative to reductionism. At this level, as popularly claimed, simulation is a new way of doing science. In addition, computer simulation can be used at a lower level, as a tool in reductionism. It is this second use of simulation that is implied by the more familiar triangular diagram shown in Figure 1.2 [15]. That diagram suggests two reductionistic uses for simulation: (a) simulation data on models can be used to test

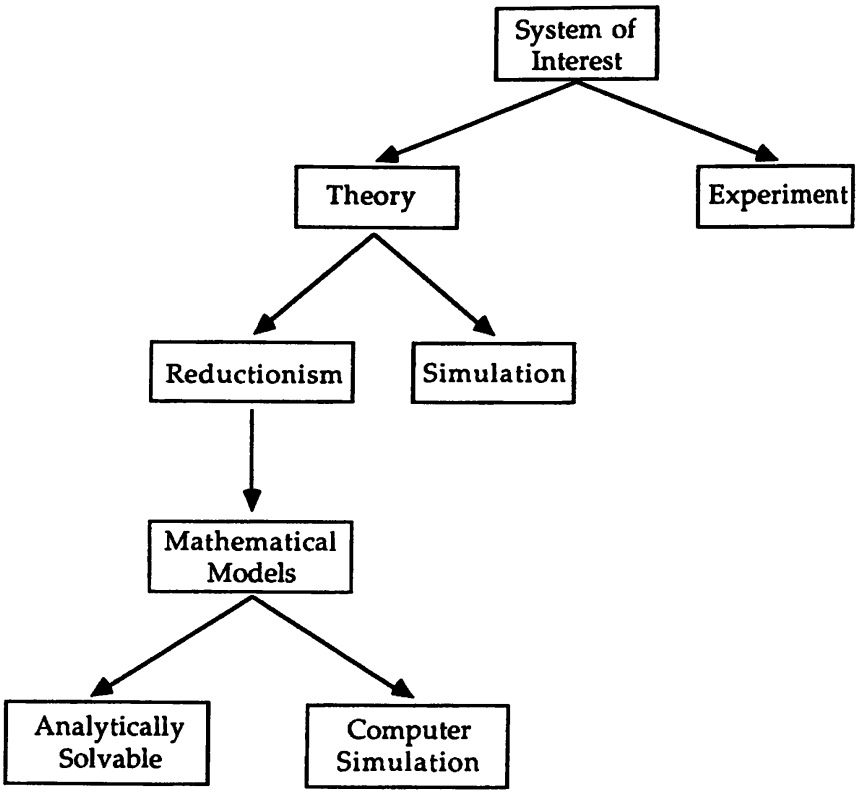

FIGURE 1.1 Hierarchy of scientific modes of investigation. Note that the system of interest may be real, or it may itself be a model.

theories and (b) simulation data can be compared with experimental data to test the realism of simulated models.

In Figure 1.1, the dual use—as an alternative to reductionism or as a reductionistic tool—explains, perhaps, early debates over whether computer simulation is theory or experiment. If you identify all theory as reductionism and sense that simulation is something different, then you may interpret simulation as experiment. Conversely, if you see computer simulation as a tool for studying reductionistic models, then you may interpret simulation as theory.

The theme of this section is that computer simulation offers possibilities more instructive and more far-reaching if it is used as an alternative to reductionism rather than as merely a servant to reductionism. To make this

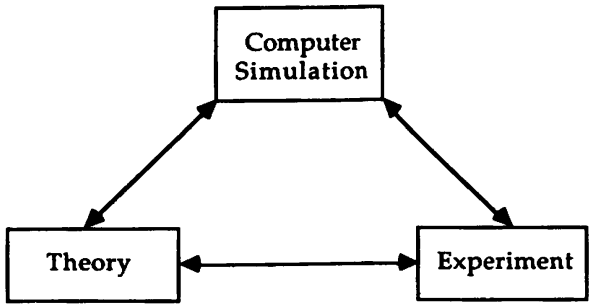

FIGURE 1.2 Conventional representation of the interplay among theory, experiment, and computer simulation.

statement concrete: using computer simulation to map out a phase diagram for a model of methane (or whatever), merely to test how realistic the model may be, is to misunderstand and underutilize the power of simulation.

1.5 MODELS FOR MOLECULAR SIMULATIONS

A computer simulation is valuable because it is applied to a precisely defined model for the material of interest. The model is actually a composite of two: one for interactions among the molecules making up the system and another for interactions between the molecules and their environment:

Simulated model	=	model for molecular interactions
	+	model for system–environment interactions

Note the decoupling implied by this schematic—intermolecular interactions are presumed to be independent of interactions with the environment.

The model for molecular interactions is contained in an intermolecular force law or, equivalently, an intermolecular potential energy function. This potential function implicitly describes the geometric shapes of individual molecules or, more precisely, their electron clouds. Thus when we specify the potential function, we establish the symmetry of the molecules, whether they are rigid or flexible, how many interaction sites occupy each molecule, and so on. A detailed characterization of intermolecular potential functions may be given analytically or numerically; in any case, a quantitative form for the potential function defines a molecular model and hence the form must be chosen before a simulation can be performed.

In this book we consider only spherically symmetric molecules (atoms). For N such atoms the intermolecular potential function is represented by $\mathscr{U}(\mathbf{r}^N)$. The notation $\mathbf{r}^N$ represents the set of vectors that locate the atomic centers of mass, $\mathbf{r}^N = \{\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3, \ldots, \mathbf{r}_N\}$. When we establish values for the set $\mathbf{r}^N$, we define the *configuration* of a system. Macroscopic properties that are averages over only the set $\mathbf{r}^N$ are called *configurational* properties.

In most simulations the intermolecular potential energy is taken to be a sum of isolated pair interactions; this assumption is called *pairwise additivity*. Hence,

$$\mathscr{U} = \sum_{i<j}\sum u(r_{ij}) \tag{1.3}$$

where $u(r_{ij})$ is a pair potential energy function whose form is known and r_{ij} is the scalar distance between molecules i and j. Since no dissipative forces