CHAPTER 1

Introduction

When a PhD student in physics picks up a textbook about elementary particles, or cosmology, or condensed matter, there is little doubt about what will be found inside the covers. There are questions, perhaps, about the level and style of presentation, or about the emphasis given to different subfields, but the overall topic is clear. The situation is very different for books or courses that attempt to bring the intellectual style of physics to bear on the phenomena of life. The problem is not just in how we teach, but also in how we do research. The community of physicists interested in biological problems is incredibly diverse, it spills over into more amorphously defined interdisciplinary communities, and individual physicists often are more connected to biologists working on the same system than they are to physicists asking the same conceptual question in other biological systems. None of this is necessarily good or bad, but it can be confusing for students.

1.1 About Our Subject

Ours is not a new subject, but over its long history, "biophysics" or "biological physics" has come to mean many different things to different communities.[1] At the same time, for many physicists today, biophysics remains new and perhaps a bit foreign. There is an excitement to working in a new field, and I hope to capture this excitement. Yet our excitement, and that of our students, sometimes is tempered by serious concerns, which can be summarized by naive questions: Where is the boundary between physics and biology? Is biophysics really physics or just the application of methods from physics to the problems of biology? My biologist friends tell me that "theoretical biology" is nonsense, so what would theoretical physicists be doing if they got interested in this field? In the interaction between physics and biology, what happens to chemistry? How much biology do I need to know to make progress? Why do physicists and biologists

1. The use of these two different words also is problematic. I think that, roughly speaking, "biophysics" can be used by people who think of themselves either as physicists or biologists, whereas "biological physics" is an attempt to carve out a subfield of physics, distinct from biology. The difficulty is that neither word really points to a set of questions that everyone can agree on. So, we need to dig in.

seem to be speaking such different languages? Can I be interested in biological problems and still be a physicist, or do I have to become a biologist? Although there has been much progress over the past decade, I still hear students (and colleagues) asking these questions, and so it seems worth a few pages to place the subject of this book into context.

To put things in perspective, we need to have in mind at least a cartoon sketch of the intellectual landscape that we are trying to explore. I am not a sociologist or historian of science, and cartoons can be, by definition, cartoonish, so there are myriad dangers here. Indeed, the dangers are sufficiently great that at least one colleague recommended scrapping this discussion completely. There is a school of thought that says, roughly, "I don't care if it's biology or physics, I just want to do interesting science," and according to this view one need not worry about the place of biophysics in relation to the larger, separate enterprises of physics and biology. This sounds good and hints at a nirvana in which disciplinary boundaries are erased and (by extension) universities have no departments. But this position is not interdisciplinary or multidisciplinary, it's antidisciplinary, and it assumes, implicitly, that the definition of "interesting" is objective, independent of our intellectual backgrounds. I find this very hard to believe.[2] One might hope that we could construct an objective definition of interesting science, but certainly practicing scientists don't automatically subscribe to such a thing. Indeed, faced with the same natural phenomena, my experience is that physicists and biologists (and mathematicians and chemists and engineers) will ask different questions. By the time we find answers, we ought to be able to convince one another that we have accomplished something. But at the stage where we are still formulating questions, I think that any search for unanimity about what is a "good" or "interesting" question would devolve into one group insisting that they alone have the right to determine what is relevant. I believe there really are physics problems motivated by the phenomena of life and that these problems can be different from those that engage our biologist friends—not better or worse, but different. So, with this in mind, let us draw that sketch, cartoonish though it may be.

Academic disciplines can define themselves either by their objects of study or by their style of inquiry. Physics is firmly in the second camp. Physicists make it their business to ask certain kinds of questions about Nature and to seek certain kinds of answers. The aspects of the world which capture the interest of the physics community can and do change, not least as new phenomena become accessible to the physicists' style of inquiry. Throughout these changes, "thinking like a physicist" is supposed to mean something, and it is this, above all else, that we try to convey to our students. We take ourselves (not without hubris) to be the intellectual heirs of Galileo, embracing his evocative claim that the book of Nature is written in the language of mathematics.

Biology translates from its Greek roots as the study of life. The style of inquiry may change, from studies of animal behavior and anatomy to genetics and molecular

2. Full disclosure: As an undergraduate at Berkeley, I took philosophy of science from Paul Feyerabend, who was a particularly eloquent (and amusing) critic of the idea that there is something objective that we could point to as a universally interesting scientific question; this was part of a larger critique of attempts to codify "the scientific method." I think his arguments were meant to make us look more carefully at our own assumptions about the nature of the scientific enterprise, and I think this worked for me. Some found him too much and construed his approach as antiscientific. I encourage you to explore for yourself in the references for this section at the end of the book.

structure, but the objects remain the same. It is especially important for physicists to appreciate the vastness of the enterprise that is labeled "biology" and the divisions within biology itself. A geneticist, for example, studying the dynamics of regulatory networks in a simple organism, such as yeast, has relatively little in common with a colleague who studies the dynamics of neural networks for the regulation of movement in higher organisms. Not only is biology defined by the objects of study, but also subfields of biology are similarly defined, so that networks of neurons and networks of genes are different subjects.

Differences in our view of the scientific enterprise translate rather directly into different educational structures. In physics, we (try to) teach principles and derive the predictions for particular examples. In biology, teaching proceeds (mostly) from example to example, system to system. Although physics has subfields, to a remarkable extent the physics community clings to the romantic notion that Physics is one subject. Not only is the book of Nature written in the language of mathematics, but also there is only one book, and we expect that if we really grasped its content, it could be summarized in very few pages. Where does biophysics fit into this view of the world?

There is something different about life, something that we recognize immediately as distinguishing the animate from the inanimate. But we no longer believe that there is a fundamental "life force" that animates a lump of inert stuff.[3] Similarly, there is no motive force that causes superfluid helium to crawl up the sides of a container and escape, or causes electrical current in a superconducting loop to flow forever; the phenomena of superfluidity and superconductivity emerge as startling consequences of well-known interactions among electrons and nuclei, interactions that usually have much more mundane consequences. As physicists studying the phenomena of life, we thus are not searching for a new force of Nature. Rather, we are trying to understand how the same forces that usually cause carbon-based materials to look like rocks or sludge can, under some conditions, cause some of this material to organize itself and walk (or swim or fly) out of the laboratory. What is special about the state of matter that we call life? How does it come to be this way? Different generations of physicists have approached these mysteries in different ways.

Looking Back

Some of the giants of classical physics—Helmholtz, Maxwell, and Rayleigh, to name a few—routinely crossed borders among disciplines that we now distinguish as physics, chemistry, biology, and even psychology. Some of their forays into the phenomena of life were driven by a desire to test the universality of physical laws, such as the conservation of energy. A very different motivation was that our own view of the world is determined by what we can see and hear, and more subtly by what we can reliably infer from the data that our eyes and ears collect. These physicists thus were drawn to the study of the senses; for them, there was no boundary between optics and vision, or between acoustics and hearing. Helmholtz in particular took a very broad view, seeing a path not just from acoustics to the mechanics of the inner ear and from the properties of light

3. This now obvious statement reflects centuries of hard work, and not a little real fighting. Einstein's explanation of Brownian motion, among other things, put to rest the idea that the spontaneous movements of microscopic particles reflected a life force.

to the optics of the eye, but all the way from the physical stimuli reaching our sense organs to the nature of our perceptions, to our ability to learn about the world, and even to what makes some sights or sounds more pleasing than others. Reading Helmholtz today I find myself struck by how much his insights still guide our thinking about vision and hearing, and how the naturalness of his cross-disciplinary discourse remains something that few modern scientists achieve, despite the current fanfare about the importance of multidisciplinary work. Most of all, I am struck by his soaring ambition that physics should not stop at the point where light hits our eyes or sound enters our ears, and that we should search for a physics that reaches all the way to our personal, conscious experience of the world in all its beauty.

The rise of modern physics motivated another wave of physicists to explore the phenomena of life. Fresh from the triumphs of quantum mechanics, they were emboldened to seek new challenges and brought new concepts. Bohr wondered aloud if the ideas of complementarity and indeterminacy would limit our ability to understand the microscopic events that provide the underpinnings of life. Delbrück was searching explicitly for new principles, hoping that a modern understanding of life would be as different from what came before as quantum mechanics was different from classical mechanics. Schrödinger, in his influential series of lectures titled *What Is Life?*, seized on the discovery that our precious genetic inheritance was stored in objects the size of single molecules, highlighting how surprising this is for a classical physicist, and contrasted the order and complexity of life with the ordering of crystals. Along the way he outlined a strikingly modern view of how nonequilibrium systems can generate structure out of disorder, continuously dissipating energy.

In one view of history, there is a direct path from Bohr, Delbrück, and Schrödinger to the emergence of molecular biology. Certainly, Delbrück did play a central role, not least because of his insistence that the community should focus (as the physics tradition teaches us) on the simplest examples of crucial biological phenomena, reproduction and the transmission of genetic information. The goal of molecular biology to reduce these phenomena to interactions among a countable set of molecules surely echoed the physicists' search for the fundamental constituents of matter, and perhaps the greatest success of molecular biology is the discovery that many of these basic molecules of life are universal, shared across organisms separated by hundreds of millions of years of evolutionary history. Where classical biology emphasized the complexity and diversity of life, the first generation of molecular biologists emphasized the simplicity and universality of life's basic mechanisms, and it is not hard to see this as an influence of the physicists who came into the field at its inception.

Another important idea at the start of molecular biology was that the structure of biological molecules matters. Although modern biology students, even in many high schools, are taught that "structure determines function," this was not always obvious. To imagine, in the years immediately after World War II, that all of classical biochemistry and genetics would be reconceptualized once we could see the actual structures of proteins and DNA was a revolutionary vision—one shared by only a handful of physicists and the most physical of chemists. Every physicist who visits the grand old Cavendish Laboratory in Cambridge should pause in the courtyard and realize that on that ground stood the Medical Research Council (MRC) hut, where Bragg nurtured a small group of young scientists who were trying to determine the structure of biological molecules through a combination of X-ray diffraction experiments and pure theory.

To make a long and glorious story short, they succeeded, perhaps even beyond Bragg's wildest dreams, and some of the most important papers of twentieth-century biology thus were written in a physics department.

Perhaps inspired by the successes of their intellectual ancestors, each subsequent generation of physicists offered a few converts. The idea, for example, that the flow of information through the nervous system might be reducible to the behavior of ion channels and receptors brought one group, armed with low-noise amplifiers, intuition about the interactions of charges with protein structure, and the theoretical tools to translate this intuition into testable quantitative predictions. The possibility of isolating a single complex of molecules that carried out the basic functions of photosynthesis brought another group, armed with the full battery of modern spectroscopic methods that had emerged in solid state physics. Understanding that the mechanical forces generated by a focused laser beam are on the same scale as the forces generated by individual biological molecules as they go about their business brought another generation of physicists to our subject. The sequencing of whole genomes, including our own, generated the sense that the phenomena of life could, at last, be explored comprehensively, and this inspired yet another group. Proximal stimuli were not always experimental: the idea that networks of neurons could be described in the language of statistical mechanics energized a whole theoretical community, and their results gradually fed back into the investigation of cells, circuits, and systems in the brain. These examples are far from complete, but they give a sense of the diversity of challenges that drew physicists toward problems that traditionally had been in the domain of biologists.

Through these many generations, some conventional views arose about the nature of science at the borders between physics and biology. First, there is a strong emphasis on technique. From X-ray diffraction to the manipulation of single molecules to functional imaging of the human brain, it certainly is true that physics has developed experimental techniques that allow much more direct exploration of questions raised by biologists. Second, there is a sense that in some larger classification system, biophysics is a biological science. Certainly when I was a student, and for many years afterward, physicists would speak (sometimes wistfully) of colleagues who were fascinated by the phenomena of life as having "become biologists." For their part, biologists would explain that physicists were successful in these explorations only to the extent that they appreciated what was "biologically important." Finally, biophysics has come to be organized along the lines of the traditional biological subfields. As a result, the dynamics of ion channels in single neurons and the collective behavior of large neural networks are separate subjects, and the generation of physicists exploring noise in the regulation of gene expression is disconnected from the previous generation that studied noise in ion channels.

Without taking anything away from what has been accomplished, I believe that much has been lost in the emergence of the conventional views about the nature of the interaction between physics and biology. By focusing on methods, we miss the fact that, faced with the same phenomena, physicists and biologists will ask different questions. In speaking of biological importance, we ignore the fact that physicists and biologists have different definitions of understanding. By organizing ourselves around structures that come from the history of biology, we lose contact with the dreams of our intellectual ancestors that the dramatic qualitative phenomena of life should be clues to deep theoretical insights, that there should be a physics of life and not just the physics of this or that particular system. It is, above all, these dreams that I want to rekindle in my students and in the readers of this book.

Looking Forward

At present, most questions about how things work in biological systems are viewed as questions that must be answered by experimental discovery. The situation in physics is very different, in that theory and experiment are more equal partners. In each area of physics we have a set of general theoretical principles, all interconnected, which define what is possible; the path to confidence in any of these principles is built on a series of beautiful quantitative experiments that have extended the envelope of what we can measure and know about the world. Beyond providing explanations for what has been seen, these principles provide a framework for exploring, sometimes playfully, what *ought* to be seen.[4] In many cases these predictions are sufficiently startling that to observe the predicted phenomena (a new particle, a new phase of matter, fluctuations in the radiation left over from the big bang, etc.) still constitutes a dramatic experimental discovery.

Can we imagine a physics of biological systems that reaches the level of predictive power that has become the standard in other areas of physics? Can we reconcile the physicists' desire for unifying theoretical principles with the obvious diversity of life's mechanisms? Could such theories engage meaningfully with the myriad experimental details of particular systems, yet still be derivable from succinct and abstract principles that transcend these details? For me, the answer to all these questions is an enthusiastic "yes." I hope that this book will convey both my enthusiasm and the reasons that lie behind it.

Although we aim at questions that have the generality and abstraction that are familiar from our experience in the rest of physics, our attention is attracted first not by

4. Reading an earlier draft, one friend found my use of "ought" here quite jarring, as if theoretical predictions were moral imperatives: children should mind their manners, and the Higgs boson should have a mass of 125 GeV. Surely physics does not make normative claims that the world should be one way or another, but rather it tries to describe the world as we find it. I was about to choose another word when I found myself reading David Foster Wallace's essay "Authority and American Usage." In what is nominally a book review of a dictionary (!), Wallace discusses many things, including the distinction between prescriptive and descriptive approaches to language. Descriptivists write dictionaries with the goal of cataloguing the language as it is used, without claiming that some patterns of usage are superior to others, only more common, whereas prescriptivists write dictionaries with the goal of defining what is correct, establishing norms for usage. According to Wallace, the descriptivists often claim their program to be the more scientific, being based on data about what native speakers actually do. One can hear echoes of this conflict in the interactions between the intellectual traditions of physics and biology. A more empirically minded biologist might insist that models be based on "everything that we know is there" in the real system. A more theoretically inclined physicist would feel comfortable writing down simplified models that leave out many details; indeed, much of the art of theorizing lies in separating what we suspect is essential from the ignorable details. Such theories thus make a normative claim, namely, that the world should behave in a way such that the things we leave out of our models aren't important. So, I think the use of "ought" in this context really gets at an essential aspect of what we do when we construct theories in the physics tradition: the theoretical physicist's relation to Nature is prescriptivist. I cannot resist noting that Wallace comes down on the side of prescriptivism, not least because a genuinely complete and unbiased descriptivism is impossible, much as it may be impossible to know everything about the components and interactions in a complex biological system. In the end, even descriptivists make normative claims about what constitutes a "good" data set, and so we might as well be honest about our prescriptions.

abstractions but by the concrete and dramatic phenomena of life, so we should start there. There are so many beautiful things about life, however, that it can be difficult to choose an appropriate starting point. Before explaining the choices I made in writing this book, I emphasize that there are many equally good choices. Indeed, if we choose almost any of life's phenomena—the development of an embryo, our appreciation of music, the ability of bacteria to live in diverse environments, the way ants find their way home in the hot desert—we can see glimpses of fundamental questions even in the seemingly most mundane events.

It is a remarkable thing that, pulling on the threads of one biological phenomenon, we can unravel so many general physics questions. In any one case, some problems will be presented in purer form than others, but in many ways everything is there. Thus, if we think hard about how crabs digest their food (to choose a particularly prosaic example), we will find ourselves worrying about how biological systems manage to find the right operating point in very large parameter spaces. This problem, as we will see in Chapter 5, arises in many different systems, across levels of organization from single protein molecules to short-term memory in the brain. Thus, in an odd way, everything is fair game. The challenge is not to find the most important or "fundamental" phenomenon but rather to see through any one of many interesting and beautiful phenomena to the deep physics problems that are hiding underneath the often formidable complexity of these systems.

The first problem, as noted above, is that there really is something different about being alive, and we'd like to know what this is—in the same way that we know what it is for a collection of atoms to be solid, for a collection of electrons to be superconducting, or for the vacuum to be confining (of quarks). This "what is life?" question harkens back to Schrödinger, and one might think that the molecular biology that arose in the decades after his manifesto would have answered his question, but this isn't clear. Looking around, we more or less immediately identify things that are alive, and the criteria that we use in making this discrimination between animate and inanimate matter surely have nothing to do with DNA or proteins. Even more strongly, we notice that things are alive long before we see them reproduce, so although self-reproduction might seem like a defining characteristic, it does not seem essential to our recognition of the living state. Being alive is a macroscopic state, whereas things like DNA and the machinery of self-reproduction are components of the microscopic mechanism by which this state is generated and maintained.[5] Although we have made much progress on identifying microscopic mechanisms, we have made rather less progress on identifying the "order parameters" that are characteristic of the macroscopic state.

Asking for the order parameters of the living state is a hard problem, and one that is not terribly well posed. One way to make progress is to realize that as we make more quantitative models of particular biological systems, these models belong to families: we can imagine a whole class of systems, with varying parameters, of which the one

5. More precisely, all the molecular components of life that we know about comprise *one way* of generating and maintaining the state that we recognize as being alive. We don't know if there are other ways, perhaps realized on other planets. This remark might once have seemed like science fiction, and perhaps it still is, but the discovery of planets orbiting distant stars has led many people to take these issues much more seriously. Designing a search for life on other planets gives us an opportunity to think more carefully about what it means to be alive.

we are studying is just one example. Presumably, most of these possible systems are not functional, living things. What then is special about the regions of parameter space that describe real biological systems? This is a more manageable question, and it can be asked at many different levels of biological organization. If there is a principle that differentiates the genuinely biological parts of parameter space from the rest, then we can elevate this principle to a theory from which the properties of the biological system could be calculated a priori, as we do in other areas of physics.

If real biological systems occupy only a small region in the parameter space of possible systems, we have to understand the dynamics by which parameters arrive at these special values. At one extreme, this is the problem of the origin of life. At the opposite extreme, we have the phenomena of physiological adaptation, whereby cells and systems adjust their behavior in relation to varying conditions or demands from the environment, sometimes in fractions of a second. In between we have learning and evolution. Adaptation, learning, and evolution represent very different mechanisms, on different but perhaps overlapping time scales, for accomplishing a common goal—tuning the parameters of a biological system to match the problems that organisms need to solve as they try to survive and reproduce. What is the character of these dynamics? Are the systems we see around us more or less "equilibrated" in these dynamics, or are today's organisms strongly constrained by the nature of the dynamics itself? Put another way, if evolution is an algorithm for finding better organisms, are the functional behaviors of modern biological systems significantly shaped by the algorithm itself, or can we say that the algorithm solves a well-defined problem, and what we see in life are the solutions to this problem?

To survive in the world, organisms do indeed have to solve a wide variety of problems. Many of these are really physics problems: converting energy from one form to another, sensing weak signals from the environment, controlling complex dynamical systems, transmitting information reliably from one place to another (or across generations), controlling the rates of thermally activated processes, predicting the trajectories of multidimensional signals, and so on. While it is obvious (now!) that everything that happens in living systems is constrained by the laws of physics, such physics problems in the life of the organism highlight these constraints and provide a special path for physics to inform our thinking about the phenomena of life.

Identifying all the physics problems that organisms need to solve is not so easy. On the one hand, thinking about how single-celled organisms, with sizes on the scale of one micron, manage to move through water, we quickly get to problems that have the look and feel of problems that we might find in Landau and Lifshitz. On the other hand, it was a truly remarkable discovery that all cells have built Maxwell demons, and that our description of a wide variety of biochemical processes can be unified by this observation (see Section 4.5). Efforts in this direction can be very rewarding, however, because they identify questions that connect functionally important behaviors—for which evolution might select—with basic physical principles. Physics shows us what is hard about these problems and where organisms face real challenges. Physics also places limits on what is possible, and this gives us an opportunity to put the performance of biological systems on an absolute scale. It makes precise our intuition that organisms are really good at solving some very difficult problems.

Let us conclude this discussion with a tentative summary. The business of life involves solving physics problems, and these problems provide us with a natural subject matter. In particular, these problems focus our attention on the concept of function, which is not part of the conventional physics vocabulary[6] but clearly is essential if we want to speak meaningfully about life. Of the possible mechanisms for solving these problems, most combinations of the available ingredients probably don't work, and specifying this functional ensemble provides a manageable approach to the larger question of what characterizes the living state. Adaptation, learning, and evolution allow organisms to find these special regions of parameter space, and the dynamics of these processes provide another natural set of problems.

1.2 About This Book

This book has its origins in a course that I have taught for several years at Princeton University. It is aimed at PhD students in physics, although a sizable number of brave undergraduates have also taken the course, as well as a handful of graduate students from biology, engineering, applied math, and other disciplines.[7] Bits and pieces have been tested in shorter courses, sometimes for quite different audiences, at the Marine Biological Laboratory, Les Houches, the Boulder Summer School on Condensed Matter Physics, "Sapienza" Università di Roma, the Rockefeller University, and the Graduate Center of the City University of New York.

Rationale

Many courses on biophysics address an audience drawn from multiple disciplines, trying both to introduce physics students to the intellectual challenges posed by the phenomena of life and to introduce biology students to the concepts and methods of physics. Doing this for graduate students is incredibly hard. In practice, making biophysics an interdisciplinary course means that the level of mathematics and physics that one draws on must be much lower than in other courses for physics graduate students. By itself this need not be a bad thing, and surely it makes the material more accessible. But if this is the only way that we teach the subject, then there is a danger that the lack of connection to deeper physics ideas will become a self-fulfilling prophecy. My plan in teaching was thus the opposite of the interdisciplinary course—to produce a biophysics course that serves as an advanced physics graduate course in the same way that my colleagues teach condensed matter physics, quantum field theory and particle physics, or astrophysics and cosmology.[8] When I was student, most physicists would have been skeptical about the feasibility of such a project, but things have changed. Physicists have explored the phenomena of life at many different levels, from single molecules to entire populations of organisms, and many have brought back insights

6. This is not quite fair. In thermodynamics we distinguish "useful work," which provides a notion of function, at least in the limited context of heat engines. But we need something more general if we want to capture the full range of problems that organisms have to solve.

7. With apologies to colleagues elsewhere in the world, I use the U.S. terms for students at different levels of their education.

8. By now this is more than a goal; it is something of a contractual obligation. Our graduate students are required to take several courses in different areas to demonstrate the breadth of their knowledge of physics, and biophysics is one of the courses on the list that can satisfy this requirement. For this plan to be meaningful, all these courses need to be taught on the same level.

that are exciting as physics, not just as applications of physics to other fields. Part of my point in teaching and in writing this book has been to celebrate this transformation.

In its earliest incarnations, the course consisted of a series of case studies—problems where physicists have tried to think about some particular biological system. The hope was that in each case study we might catch a glimpse of some deeper and more general ideas. As the course evolved, I tried to shift the balance from examples toward principles. The difficulty, of course, is that we don't know the principles—we just have candidates. At some point I decided that this was OK, and that trying to articulate the principles was important even if we get them wrong. I believe that, almost by definition, something we will recognize as a theoretical physics of biological systems will have to cut across the standard subfields of biology, organizing our understanding of very different systems as instantiations of the same underlying ideas.

Although we are searching for principles, we start by being fascinated with the *phenomena* of life. Thus, I start with one particular biological phenomenon that holds, I think, an obvious appeal for physicists: the ability of the visual system to count single photons. As we explore this phenomenon, we will meet some important facts about biological systems, see methods and concepts that have wide application, and identify and sharpen a series of questions that we can recognize as physics problems. The really beautiful measurements that people have made in this system also provide a compelling antidote to the physicists' prejudice that experiments on biological systems are necessarily messy; indeed, I think these measurements set a standard for quantitative experiments on biological systems that should be more widely appreciated and emulated.[9] Another crucial feature of the photon-counting problem is that it cuts across many levels of biological organization, from the quantum (really) dynamics of single molecules to the macroscopic dynamics of human cognition.

I think one of the most important aspects of our field is the process of digging around in the phenomena to find interesting physics problems. Obviously there is a matter of taste here, and even among physicists with similar educations different problems will leap out at us from the complexities of real biological systems. But teaching biophysics surely involves teaching this process of formulating problems, not by fiat but by going through real examples.[10] So, we will take an explicit pause to formulate problems and articulate candidate principles (Chapter 3). What emerges are three broad ideas: the importance of noise, the need for living systems to function without fine tuning of parameters, and the possibility that many different problems solved by living organisms are different aspects of one big problem about the representation of infor-

9. Perhaps surprisingly, many biologists share the expectation that their measurements will be noisy. Indeed, some biologists insist that physicists have to get used to this and that this is a fundamental difference between physics and biology. Certainly it is a difference between the sciences as they are practiced, but the claim that there is something essentially sloppy about life is deeper, and deserves more scrutiny. One not-so-hidden agenda in my course is to teach physics students that it is possible to uncover precise, quantitative facts about biological systems in the same way that we can uncover precise quantitative facts about nonbiological systems, and that this precision matters.

10. Perhaps teaching in any field should involve teaching the process of formulating problems, but in more mature areas of physics there are so many successful examples that we often feel justified in assuming that this process will become clear en passant. This may or may not be correct. One might also note that the success of theory changes this process, because one of the roles of theory is to define what would be interesting or surprising.

mation. Each of these ideas is something that many people have explored, and I hope to make clear that these ideas have generated real successes. The greatest successes, however, have been when these theoretical discussions are grounded in experiments on particular biological systems. As a result, the literature is fragmented along lines defined by the historical subfields of biology. The goal here is to present the discussion in the physics style, organized around principles from which we can derive predictions for particular examples.

My choice of candidate principles is personal, and I do not expect that everyone in the field will agree with me; indeed, by the time there is consensus, perhaps the field won't be as much fun. Even more than my choice of principles, my choice of examples is not meant to be canonical, but illustrative. In choosing these examples, I had three criteria. First, I had to understand what was going on, and of course this biases me toward cases that my friends and I have studied in the past. I apologize for this limitation and hope that I have been able to do justice at least to some fraction of the field. Second, I want to emphasize the tremendous breadth and depth of physics ideas that are relevant in thinking about the phenomena of life. Many students are given the impression, implicitly or explicitly, that to do biophysics one can get away with knowing less "real physics" than in other subfields, and I think this is a disastrous misconception. Third, if the whole program of finding principles is going to work, then it must be that a single principle really does illuminate the functioning of seemingly very different biological systems. Thus, I make a special effort to be sure that the set of examples for each principle cuts across the subfields of biology, in particular, across the great divide between molecular and cellular biology on the one hand and neurobiology on the other.[11] Ideally I would go even further, reaching toward the behavior of populations of organisms, the interactions among multiple species in natural ecologies, and so on. I can imagine how some of the same ideas and principles that I explore here connect to these larger scale phenomena, but I just don't know enough about these fields to do them justice, so I have to stop somewhere.

In trying to provide some perspective on our subject, in the previous section, I mentioned a number of now classic topics from across more than a century of interaction between physics and biology. I don't think it is right to teach by visiting these topics one after the other, for reasons which I hope are clear by now. However, it would be weird to take a whole course on biophysics and come out without having learned about these things. So I have tried to weave some of the classics into the conceptual framework of the course, perhaps sometimes in surprising places. There also are many beautiful things I have left out, and again I apologize to people who will find that I neglected matters close to their hearts. Sometimes the neglect reflects nothing more

11. I have also resisted the temptation to organize the book by physical scale, starting with molecules and ending (perhaps) with the brain. One can teach a great course in this way, but I think the principle is misleading. First, as I hope will be clear, there are actually physical principles about how biological systems function *as systems* that cut across these multiple scales. Second, ordering by scale encourages the illusion that we actually know how to build from our microscopic understanding of biological molecules "up" to the macroscopic behaviors of higher organisms. We certainly don't know how to do this, and it may even be misguided as a matter of principle—more is different, after all. Finally, starting with molecules is both historically and, if I may use the word, emotionally wrong. What intrigues us about life are the macroscopic behaviors of organisms, ourselves included, and careful quantitative investigations of these macroscopic phenomena set the stage for later molecular explorations.

than my ignorance, but I also felt strongly that everything I discuss should fit together into a larger picture and that it is almost disrespectful to give a laundry list of wonderful but undigested results. Thus, much was left unsaid.

User's Guide

As explained above, I am aiming at PhD students in physics, so I will assume that the reader has a strong physics background and is comfortable with the associated mathematical tools. Although many different areas of physics make an appearance, the most frequent references are to ideas from statistical mechanics. In practice, this is the area where at least U.S. students have the largest variance in their preparation. As a result, in places where my experience suggests that students will need help, I have not been shy to include (perhaps idiosyncratic) expositions of relevant physics topics that are not especially restricted to the biophysical context; this is, after all, a physics course. Some more technical asides are presented as an appendix, including both matters that students may have encountered in earlier courses and some opportunities for further exploration. Throughout the text, and especially in the appendix, I try very hard to avoid saying "it can be shown that"; the resulting text is longer but I hope more useful.[12]

While it is pretty clear how much physics I expect my readers to know, it might be less clear how much biology I am assuming. We can't get started without some grasp of the facts, but when we teach particle physics, we don't start by reading from the particle data book. Similarly, I won't start by reciting the "biological background." Good high school biology courses now are at the level of the introductory university courses from a generation ago. I think everybody getting a PhD knows that the instructions for building an organism are contained in DNA, even if they are unclear on exactly how these instructions are read out. So, rather than starting with a condensed pure biology course, I just plunge in, trying to wrestle with the biological phenomena that I find most exciting and stopping to explain things as needed.

I should warn the reader that I have never managed, despite a decade of good intentions, to cover all the material here in a one-semester course. Indeed, in writing the book I have been struck by those places where my lectures had "covered" a topic with a few remarks and some pointers to the literature, rather than a fully digested exposition. If this were a short book of lectures on limited topics, I might be forgiven for such casualness, but surely this book is now too long for such things. There still are pointers designed to encourage exploration, but these really should be pointing out to the unknown rather than pointing back toward incompletely explained background material.

If you can't cover everything, how should you use the book? Chapters 4, 5, and 6 present candidate principles, and each chapter discusses many examples. Some of these examples could be skipped, depending on your taste. Alternatively, you might think that the discussion of photon counting (Chapter 2) is a bit heavily biased toward the issues of noise, so you could skip Chapter 4 completely or at least use it very selectively (but don't miss kinetic proofreading, in Section 4.5). A completely different strategy would be to touch the material of every section in the book but not in too much depth,

12. Special thanks go to Phil Nelson, who noticed that an earlier draft rigorously avoided saying "it can be shown that," but suffered from a number of morally equivalent lapses. I hope I have caught them all!

compressing each section into a one-hour lecture—it is a quick pace, but possible, and still leaves time to pull in background from the appendix.

The most important practical comment about the structure of the book concerns the problems. I cannot overstate the importance of doing problems as a component of learning. One should go further, getting into the habit of calculating as one reads, checking that you understand all the steps of an argument and that things make sense when you plug in the numbers or make order-of-magnitude estimates.[13] For all these reasons, I have chosen (following Landau and Lifshitz) to embed the problems in the text, rather than relegating them to the ends of chapters. In some places the problems are small, really just reminding you to fill in some missing steps before going on to the next topic. At the opposite extreme, some problems are small research projects. Because progress in biophysics depends on intimate interaction between theory and experiment, some problems ask you to analyze real data, which can be found through http://press.princeton.edu/titles/9911.html.[14]

Finally, a few words about references. References to the original literature serve multiple functions, especially in textbooks. Most obviously, I should cite the papers that influenced my own thinking about the subject, acknowledging my intellectual debts. Because this text is based on a course for PhD students, citations also help launch the student into the current literature, marking the start of paths that can carry you well beyond digested, textbook discussions. In another direction, references point back to classic papers—papers worth reading decades after they were published, papers that can provide inspiration. Importantly, all these criteria are subjective, and it may be as important to explain the reasons for including a particular reference as to point to the reference itself. Thus, I have collected the references, with commentary, in an Annotated Bibliography at the end of the volume. Let me note that the reference list should *not* be viewed as a rough draft of the history of the subject nor as an attempt to establish objective priorities for some work over others.

---

13. In some sections I found it difficult to formulate manageable problems. I worry that this reflects poorly on my understanding. K.G. Wilson once remarked that if you really understand something, then you can do a series of calculations in which the accuracy of the prediction is in reasonable proportion to the effort of the calculation. A corollary seems to be that one can assign problems that allow students to test their own understanding.

14. In these problems, and also in problems that involve simulation rather than analytic calculations, I have chosen to phrase all the computations in MATLAB. This makes me slightly uneasy—although the language of mathematics is universal, programming languages really are not (at least not in the practical sense). Also, MATLAB is a commercial product, although there is a free alternative available at http://www.gnu.org/software/octave/. But numerical analysis of data certainly is essential for the field, and so there isn't really any way around making a choice from currently available options.