for students as well as practicing physicists. A more serious student would need to augment this text with something closer to a traditional approach to time series analysis.

An Introduction to Particle Accelerators

E. J. N. Wilson
Oxford U. Press, New York, 2001. $90.00, $45.00 paper (252 pp.). ISBN 0-19-852054-9, ISBN 0-19-850829-8 paper

In this short, descriptive "textbook" Edmund Wilson has written what he calls *An Introduction to Particle Accelerators*. The book, he explains, sets out to remedy "the lack of a simple introduction which reveals the physical principles . . . and which best matches the needs of a graduate engineer or physicist confronting the subject for the first time." He has not written a book of that description. But that should not deter casual readers from curling up with this paperback. From it they will learn a bit about the accelerators around the world, their technologies, and the physical principles used to create them.

The three parts of Wilson's book unfold into 14 chapters that touch briefly on essentially every aspect of particle accelerators, from history to future possibilities. The extensive table of contents and the space allocated to each topic reveal the character of the book—a travelogue survey of accelerator physics, technology, and applications.

After a brief history, *An Introduction to Particle Accelerators* eases into technical discussions of the transverse focusing of particle beams. Wilson discusses longitudinal dynamics, and then returns to transverse dynamics with imperfections and nonlinearities. A special section on electron beam dynamics and synchrotron radiation is followed by a quick stop at instabilities. In chapter 10 we finally arrive at acceleration in a particle accelerator, with a detour to radio-frequency (RF) technology. The tour winds down with a discussion of applications of accelerators and future research. An introduction for graduate study should rather cover half the material at twice the depth—or all the material at twice the length.

The book is written in a folksy style; Wilson places his hand on the reader's shoulder as he gives his tour. For example, "[A] particle oscillates in this focusing system like a small sphere rolling down a slightly inclined gutter with constant speed. . . ." Along this tour we are presented with equations, graphs, and pictures that serve predominantly as decorations for the text. The physics of particle beams is not so much developed as it is stated with assurance. From time to time Wilson falls into an abbreviated development of the analysis of particle motion in an accelerator system, but he usually apologizes—for treating betatron motion, for example, "in a rather rigorous way." Some readers, Wilson explains, "might find the following sections too confusing if we carry through all the terms from the rigorous theory into a study of imperfections, . . ." Wilson takes care of the reader, leading him gently through some of the complexities of the real accelerator world—imperfections and all.

The history of particle accelerators is being written every day and many of the early practitioners are alive and well. These, living historians might have versions that differ somewhat from Wilson's. The 1958 paper by Ernest Courant and Hartland Snyder (*Annals of Physics* **3**, 1, 1958), for example, was pivotal for its development of a powerful mathematical theory that could be applied to the practical design of all of today's modern accelerators and storage rings; Wilson might have emphasized this more (actually, the entire physics community should emphasize this more). He also distorts the Nicholas Christofilos story. Christofilos did not become a colleague of Courant's until after his (Christofilos's) contributions had been acknowledged and he was hired at Brookhaven in 1953.

Wilson's constant referral to the Courant–Snyder matrix and the Courant–Snyder beta function as the Twiss parameters and Twiss matrix is an incorrect attribution that permeates the field. Some years ago Frank Cole contacted Richard Twiss, who didn't understand why the parameters were named for him.

Finally, the student of physics should be somewhat careful regarding a confusion in the book about Liouville's Theorem, which expresses the incompressibility of phase space volume in any Hamiltonian system. One gets the impression that the invariance of the emittance is a consequence of this principle alone. Actually, the invariance also requires the linearity of Hamilton's equations; it is a dynamical consequence of their solution.

An Introduction to Particle Accelerators is probably not the right book for the graduate student in engineering or physics who is planning a career in the field. However, it is an easily accessible descriptive walk through the physics and technologies of particle accelerators. As such it could be a useful read for scientists who find that their research depends heavily on one of the many different types of accelerators in use around the world.

Ronald D. Ruth
Stanford Linear Accelerator Center
Stanford, California

Statistical Mechanics of Learning

A. Engel, C. Van den Broeck
Cambridge U. Press, New York, 2001. $110.00, $39.95 paper (329 pp.). ISBN 0-521-77307-5, ISBN 0-521-77479-9 paper

In recent years physicists with an interest in statistical mechanics, in their search for interesting problems, have strayed increasingly far from their traditional home area of actual physical systems. One distant field in which they have had a significant impact is learning theory. *Statistical Mechanics of Learning*, by Andreas Engel of the University of Magdeburg and Chris Van den Broeck of the Limburg University Center, summarizes the results that have been achieved. The authors have themselves been in the thick of this action, and they give an exceptionally lucid account not only of what we have learned but also of how the calculations are done.

Learning theory has a long history, dominated in its development by statisticians, computer scientists, and mathematical psychologists. The field's core problem is essentially one of statistical inference. The following simple example illustrates the main points: Suppose we have a function—a rule, or input–output relation—that is implemented by some machine. (I use "machine" in a very general sense; it could be an animal or some other natural phenomenon. All that is necessary is that its output depend on its input.) We do not know in detail how the machine actually works; all we can do is observe and measure its response to some set of inputs. The general question is, then: What can we infer about the function on the basis of this example set of input–output pairs? If we try to make a machine of our own based on these examples, how well can we expect it to imitate the original machine?

Like the mathematicians before them, statistical physicists naturally focus on some simple model systems for which one can hope to calculate something nontrivial. The ones most studied are simple computational networks having layered structures. One can prove that such a machine can compute any continuous function of its input variables with just one layer of simple computational units between input and output. So such studies have had quite general implications.

The unique contribution that statistical physicists were able to make to this work was, not surprisingly, the calculation of average properties in the "thermodynamic limit" in which both the size of the network and the number of examples are taken to infinity. This calculation complemented nicely a lot of other analyses that focused mostly on worst-case scenarios, often of finite networks.

To obtain generic knowledge one has to consider random distributions of examples, which places the problem in the realm of disordered system theory (as originally developed for alloys and polymers). Elizabeth Gardner and her coworker Bernard Derrida pioneered the application to these networks of methods from spin glass theory and thus opened the door to hundreds of subsequent investigations that have provided much new insight into learning systems. The models, methods, and results are the focus of Engel and Van den Broeck's book.

The book starts by orienting the nonexpert reader to the basic concepts in the field and then illustrates those concepts for the simplest kind of machine, the perceptron—a machine that simply computes a weighted sum of its inputs and gives a 1 or 0 output, depending on whether the sum is above or below a threshold. The authors then further develop the statistical mechanical framework, including the "replica" methods from spin glass theory, and the reader is given simple yet nontrivial examples of phase transitions. Subsequent chapters treat topics such as data clustering, the statistical dynamics of learning, the multifractal structure of the parameter space in a problem, and more complex networks. There is also a nice chapter relating the results obtained by these methods to those found by other techniques.

Given the highly technical nature of the calculations, the presentation is miraculously clear, even elegant. Although I have worked on these problems myself, I found, in reading the chapters, that I kept getting new insights. And for someone interested in applying these methods to other problems (perhaps joining in the current work on error-correcting codes and hard optimization problems, which are sketched in the final chapter), I can't think of a better place to learn the techniques. In fact, for readers with all levels of interest, I highly recommend this book as a way to learn what statistical mechanics can say about an important basic problem.

JOHN HERTZ
Nordic Institute for Theoretical Physics (NORDITA)
Copenhagen, Denmark

Quantum Optics in Phase Space

Wolfgang P. Schleich
Wiley-VCH, New York, 2001.
$90.00 (695 pp.).
ISBN 3-527-29435-X

Travel literature has its own vague rules and guidelines, flexible enough to include books that are little more than checklists of recommended museums, sights, and restaurants, along with careful descriptions of foreign climates both cultural and meteorological. They might even include intensely personal diaries that reveal the thinking of the writer as much as they reveal the geographic locale. Physicists can recognize analogs of all these types on their bookshelves, but the regions under review are not foreign countries but some part of physics or another.

Wolfgang P. Schleich's *Quantum Optics in Phase Space* is a new contribution to physics travel literature, and it deserves praise as a guidebook. The book, suitable for almost any physicist contemplating an expedition through the quantum jungle, is timely, published at the same time that many of the jungle's mysterious elements are being charted by physicists using the techniques of quantum optics. In effect, the book shows how the classical insights associated with wavepackets and phase space can be exploited to tame many of the jungle's nonclassical creatures.

Schleich's guided expedition begins in chapter 1 with an overview of the territory to come—something like a slide show on the evening before leaving base camp: highlights and snapshots (in no particular order) of some of the foundational ingredients and keynote topics of modern quantum optics. Schleich's overview covers two-level atoms, single modes of radiation, antibunching, squeezing, cavity quantum electrodynamics and one-atom masing, de Broglie optics, fluorescence line-splitting, and entanglement. Schleich has selected these snapshots well, so an accurate impression is quickly obtained of the several ways in which quantum optics bridges the quantum–classical border. One easily gets the impression, I think accurately, that quantum phenomena are both highly puzzling and best illuminated by quantum optical investigations in the vicinity of this not-so-well-defined border.

The tour continues with a chapter called simply Ante, basically another slide show and an important one for readers who want a quick refresher course on the apparatus of quantum theory, which they will encounter in later chapters. With admirable concision, Schleich defines and illustrates key elements of representation theory, the density matrix and quantum averages, the quantum harmonic oscillator, the concept of interaction Hamiltonian, and different approaches to time evolution. While doing so, he finds it easy to hand out an array of formulas that will be useful in what follows.

Schleich explains that the main journey will start with a visualization of quantum states using the Wigner function, which introduces easily the notions of state squeezing and state reconstruction. A review of the WKB (Wentzel-Kramers-Brillouin) method is then connected to the Berry phase, leading the reader to interference in phase space and the dynamical behavior of wavepackets. Here Schleich takes his time, profitably deconstructing the character of revivals and fractional revivals as an illustration of the always fruitful procedure of time-scale analysis.

Generalizations of the Wigner function, particularly the quantum phase space machinery of Roy Glauber and George Sudarshan, are important in quantum optics in connection with photon detection, so several chapters are devoted to aspects of field quantization and photon states (Fock, coherent, Schrödinger-cat, and so on). These are followed by discussions of fundamental photonic devices and techniques, including beam splitters, homodyne detection, interferometers, and photon-count statistics.

Field quantization is also needed for examination of the quantized atom–field interaction, and Schleich commendably includes such topics as the gauge principle and the validity of the dipole approximation, topics that are frequently dodged for convenience. Introduction of the two-level atom artifice is justified, and one is led naturally to one of the central themes of quantum optics, cavity QED, where one encounters the interaction between light and matter in the context of the famous Jaynes–Cummings