

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