Probably Approximately Correct

Reviewed by Marcus Feldman

Probably Approximately Correct
Leslie Valiant
Basic Books (2013), 195 pp., US$26.99
ISBN-13: 978-0465032716

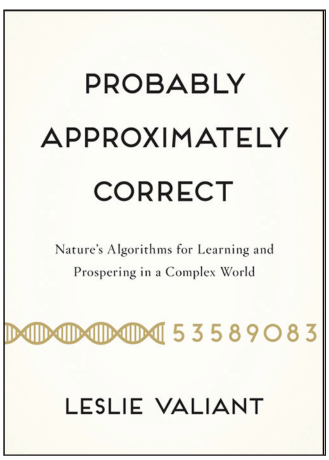

In this ambitious volume, computer scientist Leslie Valiant suggests that a vast number of natural and constructed phenomena can be described in terms of a class of algorithms he calls "ecorithms." Ecorithms are the mechanisms of processing information obtained by an entity (which could be a machine or a brain or a bacterium or a dinosaur) from its environment and using that information to improve subsequent processing in potentially different environments.

If there are good mathematical rules for predicting the process of transforming this information into knowledge, Valiant terms such processes "theoryful," while everything else is "theoryless." In both cases, whether there are precise rules for accumulating knowledge or not and whether the entity involved is a computer or an evolving species, the essence of making this transformation is learning. Theoryless processes, including evolution in biological systems or decision making in systems of cognition, are offered as innovative applications of ecorithms.

It is the computational (in the sense of computer science) features of the learning process that are characterized by the book's title. First, this learning process should take place in a relatively limited number of steps (computer scientists would say in polynomial time). Second, the number of interactions with the environment about which the entity is learning should also be limited. Third, the probability of making errors in applying the knowledge acquired by learning should be sufficiently small. A learning process that has these properties was discussed by Valiant twenty years ago and subsequently named "probably approximately correct" by D. Angluin and P. Laird in 1987, hence the title of the book and its abbreviation "PAC".

The central theme of the book is that most decisions (conscious or evolutionary) can be represented in terms of PAC learning. This learning process is described throughout the book in language that, while familiar to computer scientists, will be foreign to most biologists and cognitive scientists, who will have to work hard to (a) understand it and (b) see why it might have greater utility than the mainstream mathematical modeling in their fields.

The first two chapters expand on the above definition of PAC and quickly dive into the contributions of the great computer scientist Alan Turing to the author's thinking about ecorithms. A lot of emphasis in Chapter 3 is placed on the perceptron algorithm, a method of classifying previously unseen objects using the properties of examples called "the training set." This leads directly to the problem of induction, a problem that, in different forms, recurs throughout the book. Valiant connects induction and learning through PAC learnability, which describes the conditions on the kinds of classes that can be learned and which algorithms can be used to accomplish this learning. In this way, the heuristic concept of induction becomes quantitatively more rigorous. Both of the next chapters, "The Computable" and "The Learnable" (4 and 5), make nontrivial use of arguments from computer science, and for those who are not computer scientists, this is probably the most difficult material to absorb.

Marcus Feldman is the Burnet C. and Mildred Finley Wohlford Professor of Biological Sciences and director of the Morrison Institute for Population and Resource Studies at Stanford University. His email address is `mfeldman@stanford.edu`.

DOI: http://dx.doi.org/10.1090/noti1177

Chapters 6, 7, and 8 delve into the potential applications of PAC learning to biological evolution, human reasoning and cognition, and human emotion and language. All of these are viewed as though they involve learning processes. Evolutionists are likely to find much to disagree with in Valiant's metaphor of target pursuit for the process of biological evolution. For example, in comparing biology to physics, Valiant claims that "evolutionary theory at present offers no comparable quantitative predictions, or even quantitative explanations of the past." This fails to recognize the monumental contributions of R. A. Fisher, S. Wright, and J. B. S. Haldane, who originated the mathematical foundations of genetic evolutionary theory in the first half of the twentieth century. It also plays down the emergence of population genetic data, together with statistical tools, dynamical systems analyses, and stochastic models that were applied to explain how these patterns of genetic variation could have evolved. And finally, it ignores the development over the past twenty years of computational simulations of evolutionary trajectories that can be tested against the vast amount of genomic data that continues to accumulate.

Many population geneticists will take exception to this dismissal of the last century of work on quantitative evolutionary theory. In his goal of "Treating Darwinian evolution as a learning mechanism," Valiant implies that evolutionary biologists have been barking up the wrong evolutionary tree. In fact, he writes, "Life is full of computational mechanisms. If we are to understand how those mechanisms, and life itself, could have arisen without a designer, then computational learning is exactly where we need it to look."

In fact, the last one hundred years has seen the ascent of mathematical and statistical theory of genetic evolution to a central position in the current discourse on genomics. This theory has contributed many ideas concerning the process of speciation, quantifying fitness, and the role of modularity in evolvability, all given short shrift or not even mentioned in Chapter 6, "The Evolvable." Similarly, the fundamental work of Tversky and Kahneman in human decision making, which it seems to me would be highly relevant to Valiant's Chapter 8, "Humans as Ecorithms," is completely ignored.

While Valiant eschews the use of optimization in explicit terms, it is sometimes difficult to see why target pursuit is not optimization in PAC learnability's clothing. Explicitly, "If evolution is an instance of PAC learning, it, too, must have at least a target." Substituting "performance" for "fitness," the claim is that, in making the "target of evolution" higher performance, evolution is "amenable to treatment as a form of PAC learning" so that "exactly as in a machine learning algorithm, or ecorithm in general, the evolution algorithm will succeed without needing any expertise in, for example, ecology." Although he asserts that the "resulting creature" will not be "optimal in any sense," it is difficult to interpret evolution as learning without including what has to be learned, even if the target of learning can vary stochastically; this seems to be a general kind of approximation to optimization by evolution.

The closer the material approaches humans, the more theoryless the discussion becomes and the more tenuous the connection between PAC learning and human behavior. Valiant admits this but seems to be looking for ways to close the gap between these. There is a risk in doing this that "reductionism" can rear its ugly head. I detect some vagueness here concerning how much hard-wiring Valiant believes is required to make PAC learning a feasible explanation of such culturally variable human traits.

PAC is not an easy read. I found myself rereading many sections (I am not a computer scientist) in order to make sure that I grasped the subtle backstories (usually based in computer science) behind the development of the many sophisticated ideas in the book. It is intensely stimulating and frequently frustrating to make the connections between PAC learning, PAC semantics, and what we know about biology and decisionmaking. Besides Valiant's admitted "mentor" Turing, very few computer scientists would attempt to generate the connections among so many noncomputational disciplines that Valiant has managed in this very computational book. It is erudite, adventurous, difficult, and rewarding. It will be discussed for a long time.