


[Free pdf] The Sciences of the Artificial - 3rd Edition

The Sciences of the Artificial - 3rd Edition

Herbert A. Simon

*ePub | *DOC | audiobook | ebooks | Download PDF*



#331621 in Books imusti 1996-10-01Original language:EnglishPDF # 1 9.00 x .63 x 6.00l, #File Name: 0262691914215 pagesMit Press | File size: 40.Mb

Herbert A. Simon : The Sciences of the Artificial - 3rd Edition before purchasing it in order to gage whether or not it would be worth my time, and all praised The Sciences of the Artificial - 3rd Edition:

0 of 0 people found the following review helpful. I wish he was still alive and made many video lecturesBy Will NorskovThis book is impressive. The author's experience and awards says it all. The citation from google scholar is more than 20,000. I have read it for almost nine months and finally finished last month.It gave you so many provoking thoughts even before the big data era.5 of 8 people found the following review helpful. Review of Sciences of the

ArtificialBy June MassoudI am really enjoying reading this book. It is very well written. I understand all they are talking about because of my background in engineering, but also because I understand psychology and computer science. It's a very easy read, full of rich ideas. I really enjoy this author very much. It's not for everyone, but for people who understand the fields I mention here.27 of 30 people found the following review helpful. Heroic but flawedBy Rameez RahmanLet me first say a few words about the writing style. Simon's writing style is quite lackluster. He isn't a great writer like say Bertrand Russell, or George Orwell for that matter. But for the purposes of this book, his style suffices and is perhaps spot-on.Ok my general impression about the book: I think historically it's a groundbreaking book; it's a book written by a visionary; it's a book that at the time must have challenged a lot of people's opinions on a lot of things; in short it's an extremely important book! Having said that, here one needs to ask the all important question, reviewing it as one is, after a gap of more than 40 years since it was first published: Overall, has the book stood the test of time?The answer, surprisingly, is: `Yes' and `No'! Some of its insights are still very relevant, while some others are pretty outdated (which makes one wonder why Simon in later editions did not feel the need to say at least a few words about where he had gone wrong, and where he had over-simplified things to an astonishing degree).But before talking about both the great and not so great parts, let me briefly sketch the central idea that Simon has delineated in this book, which in fact drives the entire book. Simply, it can be described as the importance of concentrating upon the interface of a system with its outer and inner environments, without having to understand in detail either the inner or outer environments. In Simon's words, "We might look toward a science of the artificial that would depend on the relative simplicity of the interface as its primary source of abstraction and generality".Let's start with the parts that he got right. Well, first off, "Bounded rationality" of course. Simon states that the concept of bounded rationality was used by economists in some domains, even in his day (though he did coin this specific term). But he shows it quite clearly, without being antagonistic, through many examples that the concept of perfect rationality is incorrect, not only in reality (which everyone including its saner proponents accept) but also for practical purposes. Its not even good enough for practice, Simon argues persuasively. It is better to view people as bounded rational agents who adapt and satisfice rather than as perfectly rational agents who can optimize and possess an unrealistic degree of information and computational ability.Simon also says something extremely important that people tend to often forget. He highlights the fact that the debate between markets and hierarchical organizations often misses a very important empirical fact: "Roughly eighty percent of the human economic activity in the American economy, usually regarded as almost the epitome of a "market" economy, takes place in the internal environments of business and other organizations and not in the external, between-organization environments of markets". (For other, even more "shocking" facts about `free' markets, I would refer the reader to Chomsky's writings). What Simon says about organizations, about centrally planned systems utilizing markets and vice versa, is enlightening to read and merits close attention!Ok, now onto the bad parts. These are ironically the parts for which I know Herbert Simon best: Artificial Intelligence. And it is here that Simon gets it quite wrong; his vision quite flawed and again, I would say, it is quite strange that he didn't deem it fit to acknowledge his mistakes in later editions.Simon gets the abstraction terribly wrong in his ideas about the human mind. The same idea of artificial systems having simple interfaces, that works (or can work) in the sphere of human economic activity, at that particular level of abstraction, simply cannot and doesn't work when applied to the creative use of the human mind. As Simon says in this book, (and as others from the group of `cognitive revolutionaries' Chomsky, Miller et al would also say), the mind can be represented as an information processing system. But in line with the theme of the book, he feels that the mechanisms inside this information processing system are simple adaptive rules. His example of an ant finding its way back home is indicative of how extremely wrong he went with this kind of thinking. I quote:"In the case of the ant (and for that matter the others) we know the answer. He has a general sense of where home lies, but he cannot foresee all the obstacles between. He must adapt his course repeatedly to the difficulties he encounters and often detour uncrossable barriers. His horizons are very close, so that he deals with each obstacle as he comes to it; he probes for ways around or over it, without much thought for future obstacles. It is easy to trap him into deep detours. Viewed as a geometric figure, the ant's path is irregular, complex, hard to describe. But its complexity is really a complexity in the surface of the beach, not a complexity in the ant" Say what!?!Let me just state, as a matter of FACT, how ants (or other insects) do path integration is still not clear, to this very day. To make a very long point short, so that its clear to the reader where Simon gets it wrong: Even if there was no beach, even if the path to home ended up tracing a straight line, even then there would be much complexity inside the ant. The complexity of the beach pales in comparison to the complexity of what is going on inside the ant. How is it that it integrates such different cues as sun position, leg movement etc, etc, is still unknown. This is where, I feel, as far as AI is concerned, Simon's whole abstraction, his central idea, his edifice (of artificial systems) falls apart.How could a visionary like Simon go so awfully wrong? I think the answer is pretty simple. I believe it stems from the nave hopes that he (and others) had for AI, namely, simple adaptive mechanisms could give rise to "intelligence". I feel that what he said on this subject stemmed from his exuberance for AI, which can be described thus: "If the mind is an artificial system with simple adaptive rules, then we shall soon invent "intelligent" machines as well. It's only a matter of time people!"So in short, I think it was his excitement about the birth of AI, which led him to his mistakes. Like many early ambitious AI theorists, he probably

must have felt that "intelligent" machines were just around the corner. Having said that, is Simon to be blamed alone? Here's what one of my favorite scientists, the granddaddy of computing, Alan Turing says in his famous paper: "Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? If this were then subjected to an appropriate course of education one would obtain the adult brain. Presumably the child brain is something like a notebook as one buys it from the stationer's. Rather little mechanism, and lots of blank sheets. (Mechanism and writing are from our point of view almost synonymous.) Our hope is that there is so little mechanism in the child brain that something like it can be easily programmed" So we see that it is Turing who should be "blamed" for the early naive exuberance for AI. He hopes that a child's mind would be more or less a blank slate with some basic, simple mechanism. (Remember, it has to be simple so that we can program it quickly, and get our "intelligent" machine ready before Christmas). Of course now we know better. By now it's understood that even for very simple biological traits it's "fiendishly difficult," to quote a recent advanced text, to discover the genetic basis. Coming back to the book, all in all, it is a book that I would recommend highly, primarily for its historical importance, but also for the many insights that are still relevant today! In the end I would use Simon's words (and overall sentiments) that he used to describe rational action economists, to describe his own book: Heroic but flawed!

Continuing his exploration of the organization of complexity and the science of design, this new edition of Herbert Simon's classic work on artificial intelligence adds a chapter that sorts out the current themes and tools -- chaos, adaptive systems, genetic algorithms -- for analyzing complexity and complex systems. There are updates throughout the book as well. These take into account important advances in cognitive psychology and the science of design while confirming and extending the book's basic thesis: that a physical symbol system has the necessary and sufficient means for intelligent action. The chapter "Economic Reality" has also been revised to reflect a change in emphasis in Simon's thinking about the respective roles of organizations and markets in economic systems.

People sometimes ask me what they should read to find out about artificial intelligence. Herbert Simon's book *The Sciences of the Artificial* is always on the list I give them. Every page issues a challenge to conventional thinking, and the layman who digests it well will certainly understand what the field of artificial intelligence hopes to accomplish. I recommend it in the same spirit that I recommend Freud to people who ask about psychoanalysis, or Piaget to those who ask about child psychology: if you want to learn about a subject, start by reading its founding fathers. (George A. Miller *Complex Information Processing*) From the Back Cover As the Earth has made more than 5,000 rotations since *The Sciences of the Artificial* was last revised, in 1981, it is time to ask what changes in our understanding of the world call for changes in the text. About the Author Herbert Simon is Professor of Psychology at Carnegie-Mellon University. He was awarded the Nobel Prize in economics in 1978.