Book Review

Complexity: A guided tour

Melanie Mitchell (2009) New York: Oxford University Press. $29.95, 368 pages.

http://www.complexityaguidedtour.com/

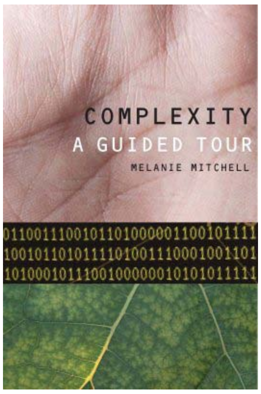

Author's information

Luis R. Izquierdo (http://luis.izquierdo.name)

Introduction

In the present day, writing a book about complexity is a major challenge. To start, there is no agreed definition of complexity. Arguably, there is not even a shared understanding of the core ideas that underlie the description of most complex systems. Concepts such as "self-organization", "adaptation", "chaos" and "emergence" are appealing and widely used, but they do not seem to have escaped the esoteric realm of ill-defined notions yet.

The challenge of writing a book on complexity becomes twofold if, in addition, the book is meant to be popular science. How can one communicate and inter-relate a myriad of complex ideas in a simple and attractive way, in fewer than 300 pages full of illustrations, using virtually no equations, and in a manner that can be easily understood by a broad audience?

Finally, if, furthermore, you want to keep a healthy skeptical viewpoint on the field and write the book in a scientifically rigorous way, the challenge becomes simply daunting.

Given all this, I honestly believe that there are very few people in the world capable of writing a decent popular science book on complexity. Some years ago I had the privilege to attend some of Prof. Mitchell's lectures at the Santa Fe Institute, and from that moment onwards I always thought that she was definitely one of those exceptional scholars who could potentially write a brilliant book on complexity. What I never suspected is that my high expectations would be exceeded. Prof. Mitchell’s book is most enjoyable, truly inspiring, giftedly written and, above all, beautifully clear.

In my opinion, the title of the book, “Complexity: A guided tour”, is particularly appropriate. The author uses the word “tour” to convey the point that complexity can be seen as a vast, culturally rich territory that contains areas that are already reasonably well tracked, but also hundreds of fascinating corners to be explored. Prof. Mitchell plays the role of the knowledgeable guide who is able to relate the history of the field with its present and its future. Importantly, the tour does not only include the most developed areas of complexity where many illuminating insights have already been reaped. Prof. Mitchell also walks the reader through less explored avenues, where scholars often differ in their views about what the core ideas are, about their significance, and about how to move the field forward. When presenting these darker areas, the author makes an outstanding effort to spell out the differences between contrasting views with exquisite clarity and scientific honesty.

Furthermore, the extensive experience of the author in conducting this type of tour for people who are unfamiliar with the field is patent throughout the book. Prof. Mitchell often seems to be able to anticipate difficulties and questions that the reader may have, and she responds to them with clarity and substantial depth. Finally, as in every tour that is meant to be understood and enjoyable for people who have never visited the territory before, no previous knowledge or specific background in math or science is required.

The remaining of this review presents the purpose of the book, a brief summary of its contents and a list of reasons why I believe the book is definitely worth reading.

Purpose of the book

The central purpose of the book is best understood by quoting the author (p. xii): “*People in the field of complex systems talk about many vague and imprecise notions such as spontaneous order, self-organization, and emergence (as well as “complexity” itself). A central purpose of this book is to provide a clearer picture of what these people are talking about and to ask whether such interdisciplinary notions and methods are likely to lead to useful science and to new ideas for addressing the most difficult problems faced by humans…*”. While widely ambitious, I do believe that the author has achieved her goal to a remarkable extent.

Summary of Contents

The book is divided into five distinct parts: “Background and History”, “Life and Evolution in Computers”, “Computation Writ Large”, “Network Thinking” and “Conclusions: The Past and Future of the Sciences of Complexity”. The following gives a brief description of each of them.

Part I: “Background and History”

The first part is the largest of the five, containing seven chapters which sum up to approximately one third of the book. It provides the history and background on the core concepts that are used throughout the rest of the book, namely: dynamics and chaos, information, computation, and evolution. Once these fundamental concepts have been presented, the first part ends with a chapter devoted to discussing and pinning down the very concept of complexity.

The first chapter starts out with a brief and informal description of five paradigmatic classes of complex systems: insect colonies, the brain, the immune system, economies, and the World Wide Web. After drawing the reader’s attention to several commonalities among these apparently disparate systems, the author presents her first introductory definition of a complex system: “*a system in which large networks of components with no central control and simple rules of operation give rise to complex collective behavior, sophisticated information processing, and adaptation via learning and evolution*”.

The author is fully aware of the fact that such a definition contains a non-negligible number of vague and imprecise terms (e.g. “simple rule”, “complex behavior”, “information processing” and “adaptation”), and the following chapters in Part I are devoted to explaining these concepts in greater depth. As a matter of fact, this first chapter ends with two statements that transpire all throughout the book:

1. “*Neither a single science of complexity nor a single complexity theory exists yet*”, and
2. “*An essential feature of forming a new science is a struggle to define its central terms*”.

The following chapters expand these two statements by giving an account of the past and the present of this quest for clarity and rigor, and by explaining how these efforts contribute to our current struggle to understand the many facets of complexity. In particular, chapter 2 -titled “Dynamics, Chaos, and Prediction”- begins by presenting the early roots of dynamical systems theory and it subsequently explains chaotic systems, the logistic map, bifurcation diagrams and Feigenbaum’s constant. The chapter ends summarizing the novelty, the essence and the relevance of chaos theory with superb clarity and brevity.

Chapter 3 tries to pin down the concept of “Information”. Naturally, it presents the view of information that is useful to the analysis of complex systems, i.e. a view where information transmission and computation does not only occur inside computers, but also within living systems. In particular, the chapter touches subjects such as statistical mechanics, entropy, and Shannon information. Unfortunately -but I guess that also inevitably- these concepts are presented in an extremely superficial way, which means that this chapter has the potential to frustrate the trained scientist.

The core subject of Chapter 4 is “Computation”. In my opinion this is one of the very best chapters of the book. Non-trivial ideas such as Gödel’s incompleteness theorem, Turing machines and the Entscheidungsproblem are explained with astonishing clarity in hardly ten pages which even include some fascinating historical notes.

Chapters 5 and 6 turn to the concept of “Evolution”. Chapter 5 gives a historical account of how the theory of evolution by natural selection was developed, starting from pre-Darwinian notions of evolution and ending with the current challenges to the Modern Synthesis. Chapter 6 is an extremely clear six-page introduction to “Genetics”. These two chapters can be perfectly recommended as a gentle, brief and lucid introduction to the history of the theory of evolution, even if the potential reader had no interest in complexity.

The final chapter of Part I is titled “Defining and Measuring Complexity”, but one could argue that a more suitable title would have been: “The Problems in Trying to Define and Measure Complexity”. The author starts by emphasizing that there are many different notions of complexity, and many of them are still very informal. She then sketches some of them (i.e. complexity as size, as Shannon entropy, as algorithmic information content, as logical depth, as thermodynamic depth, as computational capacity, as statistical complexity, as fractal dimension, and as degree of hierarchy) using the concepts presented before, and explaining both their usefulness and their theoretical and practical limitations. She finally concludes this first part of the book by stating that “*complexity probably can’t be captured by a single measurement scale*”.

Part II: “Life and evolution in Computers”

Part II contains just two chapters -“Self-reproducing Computer programs” and “Genetic Algorithms”- which are devoted to showing how life can be modeled, or mimicked, in a computer. Naturally, to do so one must start by pinning down the very concept of “life”. Thus, following her straightforward and logical approach, Prof. Mitchell starts this part of the book by outlining a few characteristics that are often considered to be necessary requisites for life, and focuses on two of them: self-reproduction and an ability to evolve or adapt.

The aim of the first of these two chapters is to build a self-reproducing computer program, i.e. a program that is able to give itself as an output. This is done in an extremely didactic and illustrative way, including a useful discussion about the dual use of information as (a) instructions to be executed and as (b) data that can be handled and operated on by those instructions. The chapter ends with a discussion about how this dual use of information relates to Gödel's incompleteness theorem and to Turing's proof of the undecidability of the Halting problem. The icing of this chapter's cake is a brief biography of Von Neumann.

The second chapter in this part illustrates how a computer program can evolve, i.e. how it can adapt itself to provide increasingly better solutions for a certain problem. This is done by explaining what genetic algorithms are, and by applying one to a particularly illustrative example (i.e. an imaginary short-sighted robot that has to clean up a 2-dimensional world containing some empty cans). The in-depth discussion of this example makes evident that genetic algorithms can often come up with useful and efficient solutions that humans had not previously considered.

Part III: "Computation Writ Large"

Having explained how life and evolution can take place inside computers, Prof. Mitchell now walks the reader through the reverse notion, i.e. how computation itself takes place within Nature. This view of Nature as a vast information processing system is getting increasingly popular in Science, and is already providing new fundamental insights on myriad disciplines (see e.g. [2]). Computation is presented here as a much broader concept than is generally understood, a concept deeply related with the meaning of life, adaptation and intelligence. In this part of the book, the author shows how this view of computation can change the way we look at living systems, can help us gain new insights about how they work, and may thus inform the design of intelligent and adaptive artificial systems.

Prof. Mitchell starts by explaining in detail why the kind of computation that takes place in Nature is fundamentally different from the kind of computation that our desktop computers with *central* processing units perform. Computation in Nature is conducted through *decentralized* systems which are most often composed of very simple units with only limited access to local information. To illustrate this fundamentally different kind of computation, the author makes use of one of the simplest models of complex systems: a cellular automaton. Cellular automata (CA) provide a wonderful way of illustrating that a collection of extremely simple deterministic components, with no central authority and with only bounded perception of their local environment, can give rise to non-trivial complex dynamics. In particular, the fact that an elementary one-dimensional CA with only two possible states is capable of universal computation [3] is especially revealing. Prof. Mitchell's introduction to CA also includes a brief discussion of Wolfram's Principle of Computational Equivalence [6].

In the following chapter the author uses an illuminating example -the "majority classification task"- to show how simple CA can perform computations that at first sight may seem to require central collective decision making. I find particularly fascinating the fact that the solution to the problem, i.e. a CA rule that is meant to successfully compute whether its initial state contains more 1s or 0s, is not simply proposed by the author; instead, the solution is evolved by means of a genetic algorithm. Thus, the example also illustrates how evolution acting upon simple units with only limited perception of their local environment can create efficient decentralized solutions to complex problems. Interestingly, the precise way in which the evolved CA works turns out to be rather difficult to understand, and is greatly clarified by the innovative use of a kind of "particle algebra" which is briefly introduced in the book. Consequently, this analysis is used to emphasize -once again- the need for a rigorous language and a set of tools to properly describe and understand complex systems.

Having shown how simple distributed systems are able to perform complex tasks in artificial worlds, Prof. Mitchell now fleshes out the extent to which one can meaningfully say that information processing or computation takes place in real decentralized complex systems. For this endeavor, she chooses the immune system, ant colonies and biological metabolism. Drawing commonalities among these systems, and using the concepts previously explained, Prof. Mitchell makes a remarkable attempt to answer four fundamental questions: "*What plays the role of "information" in these systems? How is it communicated and processed? How does this information acquire meaning? And to whom?*".

Thinking about living systems as distributed information processing networks not only helps us to better understand such living systems; it can also guide the development of biologically-inspired computer programs capable of accomplishing real-world tasks in an efficient, robust and decentralized way. These synergies have been previously exploited in e.g. the design of effective routing systems for cell phone communications using "ant colony optimization algorithms" [4], but they are particularly evident when the task to be performed has been solved successfully only by arguably the most evolved complex system in the world, i.e. our brain. This is the core of the next chapter of the book, where the author explains the functioning of Copycat, a computer program capable of drawing analogies (i.e. inferring abstract similarities between concepts in the presence of superficial differences). The venture of developing a program aimed at conducting analogies is particularly thought-provoking and illuminating since, on the one hand it forces you to put together various concepts of how distributed complex systems function, and on the other hand, once the program is written, one could say that it constitutes a formal theory about how analogical reasoning might work in our brain.

All these blue-sky computer models are surely fascinating but, are they really useful? That is the main topic of discussion for the last chapter in part III, where Prof. Mitchell takes a pause to look back and reflect on the very nature of modeling, the usefulness of simple models, and the caveats and prospects of computer modeling. This discussion is centered on the evolution of cooperation [1].

Part IV: "Network thinking"

The sophistication of most complex systems does not derive directly from the intrinsic properties of their elementary units (which are most often particularly simple), but from the patterns of interaction among them. Thus, a book about complex systems must naturally include a few chapters on the simplest way of modeling interactions formally, i.e. networks. The introduction to networks here is necessarily superficial, but certainly sufficient for the purpose of the book. Having said that, Prof. Mitchell's brief visit to the topic of networks does give her enough time to shed some refreshing light, clarity, and constructive self-critique on the field. In particular, she clearly explains that a network is scale-free if and only if it has a power-law degree distribution, she discusses the significance of power laws and the many ways one can generate them (e.g. monkeys typing randomly will do the job [5]), and she is not shy in entertaining the arguments put forward by many scientists who are skeptical about the many general claims made about power laws, their apparent and mysterious ubiquity and, more generally, about the overall significance of the so-called Network Science.

Prof. Mitchell devotes the next chapters to convincingly show how network thinking is affecting various areas of science and technology, not only by fostering novel ways to think about difficult problems, but also by providing a common language that facilitates productive cross-fertilizations among different scientific disciplines. These chapters include a discussion on Brown, Enquist and West's metabolic scaling theory (a proposed solution to one of biology's most puzzling mysteries, namely the way in which the basal metabolic rate of living organisms scales with size) and on Kauffman's work on genetic regulatory networks and their role in the evolution of complex organisms. The latter discussion is preceded by a summary of the most recent discoveries in genetics and in evolutionary developmental biology, which challenge in fundamental ways the conventional view of gradual evolution by natural selection on very small random variations in individuals.

Part V: "The Past and Future of the Sciences of Complexity"

In the final chapter of the book the author reflects on the contributions of Complex Systems research to the advancement of Science, and on the way it will probably develop in the future. I fully agree with Prof. Mitchell that the field of complexity will benefit immensely from a new, more formal, conceptual vocabulary and a new kind of mathematics -a sort of calculus of complexity, as she calls it. There's still a lot of work to do in this direction, but it is also true that there is also much useful work already done behind us, and this book is a clear testimony of that.

Conclusions

I highly recommend this book for many reasons. If you are not familiar with Complex Systems Science, here you will find an excellent introduction to the topic –definitely the clearest I have ever read. The book is beautifully written, extremely lucid and most enjoyable. Besides, a substantial part of the book is nicely structured in reasonably self-contained chapters that can be read independently, as they explain fairly distinct concepts. Thus, from now on, if someone asks me for a brief first introduction on a topic such as cellular automata, genetic algorithms, evolution, genetics, chaos, Turing Machines, fractal dimensions, scale-free networks, etc. I will recommend that this person read the relevant chapter of Prof. Mitchell's book. (Incidentally, many of these chapters correspond to introductory lectures at the Santa Fe Complex Systems Summer School, if you ever wondered what they do.) After having savored a few chapters, I am sure that reading the whole book will follow naturally; and I am glad this is the case, since -as a complex system itself- this book is more than the sum of its parts.

I would also recommend the book to someone who is already conversant with complexity. To start, given the broad range of topics and concepts covered in the book, it is likely that even experienced scientists in the field of complexity will find new ideas and perspectives they are not necessarily familiar with. A second reason to read this book, even if you are learned in complexity, comes as a positive side effect of having an incredibly broad and interdisciplinary set of concepts compressed in fewer than 300 small pages with plenty of illustrations. This constraint has forced the author to think very hard and carefully about the core of every concept she presents, distill the key issues, and explain their very quintessence in a neat, lucid and spotless style. Thus, I really enjoyed reading about topics that I was already familiar with, since I had not seen their essence explained so clearly before.

I would like to conclude this review by emphasizing the two features that, in my opinion, make this book truly outstanding: its clarity, and its healthy and balanced skepticism on various aspects of the field.

We are all familiar with the usual buzzwords in complexity: “emergence”, “self-organization”, “unpredictability”, etc. The use of these terms often leads to attractive and seemingly illuminating descriptions of complex systems, but too often these descriptions turn out to be obscure and esoteric when one looks at them with a more skeptical eye. That is certainly not the case in this book. All throughout the book, the author has made an outstanding effort to be clear and unambiguous, and to provide mathematical and empirical explanations that go beyond the less rigorous analogies characteristic of earlier attempts.

It is tempting, potentially self-satisfying and sometimes even comforting, to follow the crowd and sell complexity as the unifying field that is already providing all the answers that reductionism could not find. Tempting as it may be, the truth is that we are not quite there yet. The field is at its infancy; we have some success stories, and prospects are certainly most promising, but there is still a long way to walk ahead of us. There are areas in the field that have already provided ground-breaking insights -new ways of conceptualizing complex problems-, and have even yielded tremendously useful innovations, but there are also many dark places still to be illuminated. In my opinion, the most promising way to take the field forward in a scientifically rigorous way is to overtly recognize that there are still many core ideas that are vague and ambiguous; these ideas are often used in different ways by different researchers, so the first thing we have to do to make them more formal and precise is to spell out the differences between contrasting views. This book takes a major step forward in promoting this philosophy, valuing and encouraging constructive skepticism on the field. In the book, Prof. Mitchell openly welcomes and publicizes several critiques, and responds to them in a balanced, sensible, carefully thought, objective and rigorous way.

In short, the author’s enthusiasm and passion for the field makes the book fascinating to read. Her rigor, clarity and healthy skepticism make the book sound and the field scientifically stronger. Prof. Mitchell’s book is an excellent, fascinating, and rigorous account of the scientific field of complexity. She proves by example that it is possible to explain Complex Systems Science with rigor, breadth, depth and, above all, with exquisite clarity.

References

1. Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.
2. Chaitin, J. G. (2007). *Thinking about Gödel and Turing: Essays on complexity, 1970-2007*. Singapore: World Scientific Publishing.
3. Cook, M. (2004). Universality in elementary cellular automata. *Complex Systems*, 15 (1), 1-40.
4. Dorigo, M., & Stützle, T. (2004). *Ant Colony Optimization*. Cambridge, MA: MIT Press.
5. Miller, G. A. (1957). Some effects on intermittent silence. *The American Journal of Psychology*, 70 (2), 311-314.
6. Wolfram, S. (2002). *A new kind of science*. Champaign, IL: Wolfram Media.