This is page i
Printer: Opaque this

Algorithmic Randomness and Complexity

Rodney G. Downey
School of Mathematics, Statistics and Operations Research
Victoria University
PO Box 600
Wellington
New Zealand

Denis R. Hirschfeldt
Department of Mathematics
University of Chicago
Chicago IL 60637
U.S.A.

Rod dedicates this book to his wife Kristin, and Denis to Larisa.

Contents

Preface

Though we did not know it at the time, this book's genesis began with the arrival of Cris Calude in New Zealand. Cris has always had an intense interest in algorithmic information theory. The event that led to much of the recent research presented here was the articulation by Cris of a seemingly innocuous question. This question goes back to Solovay's legendary manuscript [371], and Downey learned of it during a visit made to Victoria University in early 2000 by Richard Coles, who was then a postdoctoral fellow with Calude at Auckland University. In effect, the question was whether the Solovay degrees of left-computably enumerable reals are dense.

At the time, neither of us knew much about Kolmogorov complexity, but we had a distinct interest in it after Lance Fortnow's illuminating lectures [148] at Kaikoura[1] in January 2000. After thinking about Calude's question for a while, and eventually solving it together with André Nies [116], we began to realize that there was a huge and remarkably fascinating area of research, whose potential was largely untapped, lying at the intersection of computability theory and the theory of algorithmic randomness.

We also found that, while there is a truly classic text on Kolmogorov complexity, namely Li and Vitányi [248], most of the questions we were in-

[1]Kaikoura was the setting for a wonderful meeting on computational complexity. There is a set of lecture notes [112] resulting from this meeting, aimed at graduate students. Kaikoura is on the east coast of the South Island of New Zealand, and is famous for its beauty and for tourist activities such as whale watching and dolphin, seal, and shark swimming. The name "Kaikoura" is a Maori word meaning "eat crayfish", which is a fine piece of advice.

terested in either were open, were exercises in Li and Vitányi with difficulty ratings of about 40-something (out of 50), or necessitated an archaeological dig into the depths of a literature with few standards in notation[2] and terminology, marked by relentless rediscovery of theorems and a significant amount of unpublished material. Particularly noteworthy among the unpublished material was the aforementioned set of notes by Solovay [371], which contained absolutely fundamental results about Kolmogorov complexity in general, and about initial segment complexity of sets in particular. As our interests broadened, we also became aware of important results from Stuart Kurtz' PhD dissertation [228], which, like most of Solovay's results, seemed unlikely ever to be published in a journal. Meanwhile, a large number of other authors started to make great strides in our understanding of algorithmic randomness.

Thus, we decided to try to organize results on the relationship between algorithmic randomness and computability theory into a coherent book. We were especially thankful for Solovay's permission to present, in most cases for the first time, the details from his unpublished notes.[3] We were encouraged by the support of Springer in this enterprise.

Naturally, this project has conformed to Hofstadter's Law: It always takes longer than you expect, even when you take into account Hofstadter's Law. Part of the reason for this delay is that a large contingent of gifted researchers continued to relentlessly prove theorems that made it necessary to rewrite large sections of the book.[4] We think it is safe to say that the study of algorithmic randomness and dimension is now one of the most active areas of research in mathematical logic. Even in a book this size, much has necessarily been left out. To those who feel slighted by these omissions, or by inaccuracies in attribution caused by our necessarily imperfect historical knowledge, we apologize in advance, and issue a heartfelt invitation to write their own books. Any who might feel inclined to thank us will find a suggestion for an appropriate gift on page 517.

This is *not* a basic text on Kolmogorov complexity. We concentrate on the Kolmogorov complexity of sets (i.e., infinite sequences) and cover only as much as we need on the complexity of finite strings. There is quite a lot of background material in computability theory needed for some of the more sophisticated proofs we present, so we do give a full but, by necessity, rapid refresher course in basic "advanced" computability theory. This material

[2]We hope to help standardize notation. In particular, we have fixed upon the notation for Kolmogorov complexity used by Li and Vitányi: C for plain Kolmogorov complexity and K for prefix-free Kolmogorov complexity.

[3]Of course, Li and Vitányi used Solovay's notes extensively, mostly in the exercises and for quoting results.

[4]It is an unfortunate consequence of working on a book that attempts to cover a significant portion of a rapidly expanding area of research that one begins to hate one's most productive colleagues a little.

should not be read from beginning to end. Rather, the reader should dip into Chapter 2 as the need arises. For a fuller introduction, see for instance Rogers [334], Soare [366], Odifreddi [310, 311], or Cooper [79].

We will mostly avoid historical comments, particularly about events predating our entry into this area of research. The history of the evolution of Kolmogorov complexity and related topics can make certain people rather agitated, and we feel neither competent nor masochistic enough to enter the fray. What seems clear is that, at some stage, time was ripe for the evolution of the ideas needed for Kolmogorov complexity. There is no doubt that many of the basic ideas were implicit in Solomonoff [369], and that many of the fundamental results are due to Kolmogorov [211]. The measure-theoretic approach was pioneered by Martin-Löf [259]. Many key results were established by Levin in works such as [241, 242, 243, 425] and by Schnorr [348, 349, 350], particularly those using the measure of domains to avoid the problems of plain complexity in addressing the initial segment complexity of sets. It is but a short step from there to prefix-free complexity (and discrete semimeasures), first articulated by Levin [243] and Chaitin [58]. Schnorr's penetrating ideas, only some of which are available in their original form in English, are behind much modern work in computational complexity, as well as Lutz' approach to effective Hausdorff dimension in [252, 254], which is based on martingales and orders. As has often been the case in this area, however, Lutz developed his material without being too aware of Schnorr's work, and was apparently the first to explicitly connect orders and Hausdorff dimension. From yet another perspective, martingales, or rather supermartingales, are essentially the same as continuous semimeasures, and again we see the penetrating insight of Levin (see [425]).

We are particularly pleased to present the results of Kurtz and Solovay mentioned above, as well as hitherto unpublished material from Steve Kautz' dissertation [200] and the fundamental work of Antonin Kučera. Kučera was a real pioneer in connecting computability and randomness, and we believe that it is only recently that the community has really appreciated his deep intuition.

Algorithmic randomness is a highly active field, and still has many fascinating open questions and unexplored directions of research. Recent lists of open questions include Miller and Nies [278] and the problem list [2] arising from a workshop organized by Hirschfeldt and Miller at the American Institute of Mathematics in 2006. Several of the questions on these lists have already been solved, however, with many of the solutions appearing in this book. We will mention a number of open questions below, some specific, some more open ended. The pages on which these occur are listed in the index under the heading *open question*.

Acknowledgments

Our foremost acknowledgment is to the Marsden Fund of New Zealand. There is no question that this work would not have happened without its generous support. Hirschfeldt was Downey's Marsden Postdoctoral Fellow when the work that eventually led to this book began. George Barmpalias, Laurent Bienvenu, Richard Coles, Noam Greenberg (now a faculty member at VUW), Evan Griffiths, Geoff LaForte, Joe Miller, and Antonio Montalbán have also been Downey's Marsden Postdoctoral Fellows. Keng Meng Ng, Adam Day, and Andrew Fitzgerald were supported by the Marsden Fund as Downey's PhD students. Downey has also received direct Marsden support for the period of the writing of this book, as has André Nies, whose work appears throughout the text. As mentioned in the preface, our early interest in Kolmogorov complexity was stimulated by a talk given by Lance Fortnow at a conference in Kaikoura almost completely supported by the Marsden Fund, via the New Zealand Mathematical Research Institute.[5] Downey and Liang Yu have also been supported by the New Zealand

[5]This institute is a virtual one, and was the brainchild of Vaughan Jones. After receiving the Fields Medal, among other accolades, Vaughan has devoted his substantial influence to bettering New Zealand mathematics. The visionary NZMRI was founded with this worthy goal in mind. It runs annual workshops at picturesque locations, each devoted to a specific area of mathematics. These involve lecture series by overseas experts aimed at graduate students, and are fully funded for New Zealand attendees. The 2009 workshop, held in Napier, was devoted to algorithmic information theory, computability, and complexity, and Hirschfeldt was one of the speakers. The NZMRI is chaired by Vaughan Jones, and has as its other directors Downey and the uniformly excellent Marston Conder, David Gauld, and Gaven Martin.

Institute for Mathematics and its Applications, a recent CoRE (Centre of Research Excellence) that grew from the NZMRI, with Yu being Downey's postdoctoral fellow supported by the logic and computability programme of this CoRE. As a postdoctoral fellow, Guohua Wu was supported by the New Zealand Foundation for Research Science and Technology, having previously been supported by the Marsden Fund as Downey's PhD student. Stephanie Reid received similar support from the Marsden Fund for her MSc thesis. Finally, many visitors and temporary fellows at VUW have been supported by the Marsden Fund, and sometimes the ISAT Linkages programme, including Eric Allender, Veronica Becher, Peter Cholak, Barbara Csima, Carl Jockusch, Steffen Lempp, Andy Lewis, Jan Reimann, Ted Slaman, Sebastiaan Terwijn, and Rebecca Weber, among others.

Downey received a Maclaurin Fellowship from the NZIMA as part of the Logic and Computation programme, and a James Cook Fellowship from the New Zealand government for 2008–2010, both of which were important in supporting his work on this book.

We also owe substantial thanks to the National Science Foundation of the United States. Hirschfeldt's research has been supported by NSF grants throughout the period of the writing of this book, including a Focused Research Group grant on algorithmic randomness involving researchers from around a dozen institutions in the United States, who together have produced a significant portion of the research presented here. Among many activities, this FRG grant has funded important workshops in the area, in Chicago in 2008, Madison in 2009, and South Bend in 2010. The group of researchers involved in this project first came together at a workshop organized by Hirschfeldt and Joe Miller at the American Institute of Mathematics. The NSF has also funded many visitors to the University of Chicago, including Downey, who has also been funded by the NSF for visits to other institutions in the United States.

Many colleagues have helped us by sending in corrections and suggestions, by helping us with various aspects of the production of this book, and by generously sharing their results and ideas. Although errors and infelicities no doubt remain, and are entirely our responsibility, the text has benefited enormously from their assistance. With apologies to those we may have forgotten, we thank Eric Allender, Klaus Ambos-Spies, Bernie Anderson, Uri Andrews, Asat Arslanov, George Barmpalias, Veronica Becher, Laurent Bienvenu, Paul Brodhead, Cris Calude, Douglas Cenzer, Gregory Chaitin, Peter Cholak, Chi Tat Chong, Richard Coles, Chris Conidis, Barry Cooper, Barbara Csima, Adam Day, Ding Decheng, Jean-Paul Delahaye, David Diamondstone, Damir Dzhafarov, Santiago Figueira, Andrew Fitzgerald, Lance Fortnow, Johanna Franklin, Peter Gács, Noam Greenberg, Evan Griffiths, John Hitchcock, Carl Jockusch, Asher Kach, Bakh Khoussainov, Bjørn Kjos-Hanssen, Julia Knight, Antonin Kučera, Geoff LaForte, Steffen Lempp, Leonid Levin, Jack Lutz, Elvira Mayordomo, Wolfgang Merkle, Joe Mileti, Joe Miller, Antonio Montalbán, André Nies, Keng Meng Ng, Chris Porter, Alex Raichev, Jan Reimann, Jeffrey Remmel, Alexander Shen, Richard Shore, Ted Slaman, Bob Soare, Robert Solovay, Frank Stephan, Jonathan Stephenson, Dan Turetsky, Nikolai Vereshchagin, Paul Vitányi, Rebecca Weber, Guohua Wu, Cathryn Young, Yang Yue, Liang Yu, Xizhong Zheng, and Marius Zimand.

We also thank the staff at Springer for all their help with the preparation of this book, and for their patience during this rather lengthy process. Particular thanks go to Ronan Nugent and Vaishali Damle. We are proud to inaugurate the CiE-Springer book series *Theory and Applications of Computability*, and thank the Computability in Europe Association, and particularly series editors Elvira Mayordomo and Barry Cooper. Finally, we thank the anonymous reviewers for their many valuable suggestions.

Introduction

What does it mean to say that an individual mathematical object such as an infinite binary sequence is random? Or to say that one sequence is more random than another? These are the most basic questions motivating the work we describe in this book. Once we have reasonable tools for measuring the randomness of infinite sequences, however, other questions present themselves: If we divide our sequences into equivalence classes of sequences of the same "degree of randomness", what does the resulting structure look like? How do various possible notions of randomness relate to each other and to the measures of complexity used in computability theory and algorithmic information theory, and to what uses can they be put? Should it be the case that high levels of randomness mean high levels of complexity or computational power, or low ones? Should the structures of computability theory such as Turing degrees and computably enumerable sets have anything to do with randomness? The material in this book arises from questions such as these. Much of it can be thought of as exploring the relationships between three fundamental concepts: relative computability, as measured by notions such as Turing reducibility; information content, as measured by notions such as Kolmogorov complexity; and randomness of individual objects, as first successfully defined by Martin-Löf (but prefigured by others, dating back at least to the work of von Mises). While some fundamental questions remain open, we now have a reasonable insight into many of the above questions, and the resulting body of work contains a number of beautiful and rather deep theorems.

When considering sequences such as

$$101010101010101010101010101010101010\ldots$$

and

$$1011010111010101111000010101000101111\ldots,$$

none but the most contrarian among us would deny that the second (obtained by the first author by tossing a coin) is more random than the first. However, in measure-theoretic terms, they are both equally likely. Furthermore, what are we to make in this context of, say, the sequence obtained by taking our first sequence, tossing a coin for each bit in the sequence, and, if the coin comes up heads, replacing that bit by the corresponding one in the second sequence? There are deep and fundamental questions involved in trying to understand why some sequences should count as "random" and others as "lawful", and how we can transform our intuitions about these concepts into meaningful mathematical notions.

The roots of the study of algorithmic randomness go back to the work of Richard von Mises in the early 20th century. In his remarkable paper [402], he argued that a sequence should count as random if all "reasonable" infinite subsequences satisfy the law of large numbers (i.e., have the same proportion of 0's as 1's in the limit). This behavior is certainly to be expected of any intuitively random sequence. A sequence such as $1010101010\ldots$ should not count as random because, although it itself satisfies the law of large numbers, it contains easily described subsequences that do not. Von Mises wrote of "acceptable selection rules" for subsequences. Wald [403, 404] later showed that for any countable collection of selection rules, there are sequences that are random in the sense of von Mises, but at the time it was unclear exactly what types of selection rules should be acceptable. There seemed to von Mises to be no canonical choice.

Later, with the development of computability theory and the introduction of generally accepted precise mathematical definitions of the notions of algorithm and computable function, Church [71] made the first explicit connection between computability theory and randomness by suggesting that a selection rule be considered acceptable iff it is computable. In a sense, this definition of what we now call *Church stochasticity* can be seen as the birth of the theory of *algorithmic* randomness. A blow to the von Mises program was dealt by Ville [401], who showed that for any countable collection of selection rules, there is a sequence that is random in the sense of von Mises but has properties that make it clearly nonrandom. (In Ville's example, the ratio of 0's to 1's in the first n bits of the sequence is at least 1 for all n. If we flip a fair coin, we certainly expect the ratio of heads to tails not only to tend to 1, but also to be sometimes slightly larger and sometimes slightly smaller than 1.)

One might try to get around this problem by adding further specific statistical laws to the law of large numbers in the definition of randomness, but there then seems to be no reason not to expect Ville-like results from reappearing in this modified context. Just because a sequence respects laws A, B, and C, why should we expect it to respect law D? And law D may be one we have overlooked, perhaps one that is complicated to state, but clearly should be respected by any intuitively random sequence. Thus Ville's Theorem caused the situation to revert basically to what it had been before Church's work: intuitively, a sequence should be random if it passes all "reasonable" statistical tests, but how do we make this notion precise? Once again, the answer involved computability theory. In a sweeping generalization, Martin-Löf [259] noted that the particular statistical tests that had been considered (the law of large numbers, the law of iterated logarithms, etc.) were special cases of a general abstract notion of statistical test based on the notion of an "effectively null" set. He then defined a notion of randomness based on passing *all* such tests (or equivalently, not being in any effectively null set).

Martin-Löf's definition turned out to be not only foundationally well-motivated but mathematically robust and productive. Now known as 1-*randomness* or *Martin-Löf randomness*, it will be the central notion of randomness in this book, but not the only one. There are certainly other reasonable choices for what counts as "effectively null" than the one taken by Martin-Löf, and many notions of randomness resulting from these choices will be featured here. Furthermore, one of the most attractive features of the notion of 1-randomness is that it can be arrived at from several other approaches, such as the idea that random sequences should be incompressible, and the idea that random sequences should be unpredictable (which was already present in the original motivation behind von Mises' definition). These approaches lead to equivalent definitions of 1-randomness in terms of, for example, prefix-free Kolmogorov complexity and computably enumerable martingales (concepts we will define and discuss in this book). Like Martin-Löf's original definition, these alternative definitions admit variations, again leading to other reasonable notions of algorithmic randomness that we will discuss. The evolution and clarification of many of these notions of randomness is carefully discussed in Michiel van Lambalgen's dissertation [397].

The first five chapters of this book introduce background material that we use throughout the rest of the text, but also present some important related results that are not quite as central to our main topics. Chapter 1 briefly covers basic notation, conventions, and terminology, and introduces a small amount of measure theory. Chapter 2 is a whirlwind tour of computability theory. It assumes nothing but a basic familiarity with a formalism such as Turing machines, at about the level of a first course on "theory of computation", but is certainly not designed to replace dedicated texts such as Soare [366] or Odifreddi [310, 311]. Nevertheless, quite sophisticated computability-theoretic methods have found themselves into the study of algorithmic randomness, so there is a fair amount of material in that chapter. It is probably best thought of as a reference for the rest of the text, possibly only to be scanned at a first reading.

Chapter 3 is an introduction to Kolmogorov complexity, focused on those parts of the theory that will be most useful in the rest of the book. We include proofs of basic results such as counting theorems, symmetry of information, and the Coding Theorem, among others. (A much more general reference is Li and Vitányi [248].) As mentioned above, Martin-Löf's measure-theoretic approach to randomness is not the only one. It can be thought of as arising from the idea that random objects should be "typical". As already touched upon above, two other major approaches we will discuss are through "unpredictability" and "incompressibility". The latter is perhaps the least obvious of the three, and also perhaps the most modern. Nowadays, with file compression a concept well known to many users of computers and other devices involving electronic storage or transmission, it is perhaps not so strange to characterize randomness via incompressibility, but it seems clear that typicality and unpredictability are even more intuitive properties of randomness. Nevertheless, the incompressibility approach had its foundations laid at roughly the same time as Martin-Löf's work, by Kolmogorov [211], and in a sense even earlier by Solomonoff [369], although its application to infinite sequences had to wait a while.

Roughly speaking, the Kolmogorov complexity of a finite string is the length of its shortest description. To formalize this notion, we use universal machines, thought of as "optimal description systems". We then get a good notion of randomness for finite strings: a string σ is random iff the Kolmogorov complexity of σ is no shorter than the length of σ (which we can think of as saying that σ is its own best description). Turning to infinite sequences, however, we have a problem. As we will see in Theorem 3.1.4, there is no infinite sequence all of whose initial segments are incompressible. We can get around this problem by introducing a different notion of Kolmogorov complexity, which is based on machines whose domains are antichains, and corresponds to the idea that if a string τ describes a string σ, then this description should be encoded entirely in the bits of τ, not in its length. In Section 3.5, we further discuss how this notion of *prefix-free* complexity can be seen as capturing the intuitive meaning of Kolmogorov complexity, arguably better than the original definition, and will briefly discuss its history. In Chapter 6 we use it to give a definition of randomness for infinite sequences equivalent to Martin-Löf's.

In Chapter 4, we present for the first time in published form the details of Solovay's remarkable results relating plain and prefix-free Kolmogorov complexity, and related results by Muchnik and Miller. We also present Gács' separation of two other notions of complexity introduced in Chapter 3, a surprisingly difficult result, and a significant extension of that result by Day. Most of the material in this chapter, although quite interesting in itself, will not be used in the rest of the book.

In Chapter 5 we discuss effective real numbers, in particular the left computably enumerable reals, which are those reals that can be computably approximated from below. These reals play a similar role in the theory of algorithmic randomness as the computably enumerable sets in classical computability theory. Many of the central objects in this book (Martin-Löf tests, Kolmogorov complexity, martingales, etc.) have naturally associated left-c.e. reals. A classic example is Chaitin's Ω, which is the measure of the domain of a universal prefix-free machine, and is the canonical example of a specific 1-random real (though, as we will see, it is in many ways not a "typical" 1-random real).

The next three chapters introduce most of the notions of algorithmic randomness we will study and examine their connections to computability theory.

Chapter 6 is dedicated to 1-randomness (and its natural generalization, n-randomness). We introduce the concepts of Martin-Löf test and of martingale, using them, as well as Kolmogorov complexity, to give definitions of randomness in the spirit of the three approaches mentioned above, and prove Schnorr's fundamental theorems that these definitions are equivalent. We also include some fascinating theorems by Miller, Yu, Nies, Stephan, and Terwijn on the relationship between plain Kolmogorov complexity and randomness. These include plain complexity characterizations of both 1-randomness and 2-randomness. We prove Ville's Theorem mentioned above, and introduce some of the most important tools in the study of 1-randomness and its higher level versions, including van Lambalgen's Theorem, effective 0-1 laws, and the Ample Excess Lemma of Miller and Yu. In the last section of this chapter, we briefly examine randomness relative to measures other than the uniform measure, a topic we return to, again briefly, in Chapter 8.

In Chapter 7 we introduce other notions of randomness, mostly based on variations on Martin-Löf's approach and on considering martingales with different levels of effectiveness. Several of these notions were originally motivated by what is now known as *Schnorr's critique*. Schnorr argued that 1-randomness is essentially a computably enumerable, rather than computable, notion and therefore too strong to capture the intuitive notion of randomness relative to "computable tests". We study various notions, including Schnorr randomness and computable randomness, introduced by Schnorr, and weak n-randomness, introduced by Kurtz. We discuss test set, martingale, and machine characterizations of these notions. We also return to the roots of the subject to discuss stochasticity, and study nonmonotonic randomness, which leads us to one of the most basic major open questions in the area: whether nonmonotonic randomness is strictly weaker than 1-randomness.

Chapter 8 is devoted to the interactions between randomness and computability. Highlights include the Kučera-Gács Theorem that any set can be coded into a 1-random set, and hence all degrees above $\mathbf{0}'$ contain 1-random sets; Demuth's Theorem relating 1-randomness and truth table reducibility; Stephan's dichotomy theorem relating 1-randomness and PA degrees; the result by Barmpalias, Lewis, and Ng that each PA degree is the join of two 1-random degrees; and Stillwell's Theorem that the "almost all" theory of the Turing degrees is decidable. We also examine how the ability to compute a fixed-point free function relates to initial segment complexity, discuss jump inversion for 1-random sets, and study the relationship between n-randomness, weak n-randomness, and genericity, among other topics. In addition, we examine the relationship between computational power and separating notions of randomness. For example, we prove the remarkable result of Nies, Stephan, and Terwijn that a degree contains a set that is Schnorr random but not computably random, or one that is computably random but not 1-random, iff it is high. We finish this chapter with Kurtz' results, hitherto available only in his dissertation, on "almost all" properties of the degrees, such as the fact that almost every set is computably enumerable in and above some other set, and versions of some of these results by Kautz (again previously unpublished) converting "for almost all sets" to "for all 2-random sets".

The next five chapters examine notions of relative randomness: What does it mean to say that one sequence is more random than another? Can we make precise the intuition that if we, say, replace the even bits of a 1-random sequence by 0's, the resulting sequence is "$\frac{1}{2}$-random"?

In Chapter 9 we study reducibilities that act as measures of relative randomness, focusing in particular, though not exclusively, on left-c.e. reals. For instance, we prove the result due to Kučera and Slaman that the 1-random left-c.e. reals are exactly the ones that are complete for a strong notion known as *Solovay reducibility*. These reals are also exactly the ones equal to Ω for some choice of universal prefix-free machine, so this result can be seen as an analog to the basic computability-theoretic result that all versions of the halting problem are essentially the same. We also prove that for a large class of reducibilities, the resulting degree structure on left-c.e. reals is dense and has other interesting properties, and discuss a natural but flawed strengthening of weak truth table reducibility known as *cl-reducibility*, and a better-behaved variation on it known as *rK-reducibility*.

In Chapter 10 we focus on reducibilities appropriate for studying the relative randomness of sets already known to be 1-random. We study K- and C-reducibility, basic measures of relative initial segment complexity introduced in the previous chapter, including results by Solovay on the initial segment complexity of 1-random sets, and later echoes of this work, as well as theorems on the structure of the K- and C-degrees. We introduce *van Lambalgen reducibility* and the closely related notion of *LR-reducibility*, and study their basic properties, established by Miller and Yu. We see that vL-reducibility is an excellent tool for studying the relative randomness of 1-random sets. It can be used to prove results about K- and C-reducibilities for which direct proofs seem difficult, and to establish theorems that help make precise the intuition that randomness should be antithetical to computational power. For instance, we show that if $A \leqslant_{\mathrm{T}} B$ are both 1-random, then if B is n-random, so is A. Turning to LR-reducibility (which agrees with vL-reducibility on the 1-random sets but not elsewhere), we introduce an important characterization due to Kjos-Hanssen, discuss structural results due to Barmpalias and others, and prove the equivalence between LR-reducibility and *LK-reducibility*, a result by Kjos-Hanssen, Miller, and Solomon related to the lowness notions mentioned in the following paragraph. We finish the chapter with a discussion of the quite interesting concept of *almost everywhere domination*, which arose in the context of the reverse mathematics of measure theory and turned out to have deep connections with algorithmic randomness.

Chapter 11 is devoted to one of the most important developments in recent work on algorithmic randomness: the realization that there is a class of "randomness-theoretically weak" sets that is as robust and mathematically interesting as the class of 1-random sets. A set A is *K-trivial* if its initial segments have the lowest possible prefix-free complexity (that is, the first n bits of A are no more difficult to describe than the number n itself). It is *low for 1-randomness* if every 1-random set is 1-random relative to A. It is *low for K* if the prefix-free Kolmogorov complexity of any string relative to A is the same as its unrelativized complexity, up to an additive constant. We show that there are noncomputable sets with these properties, and prove Nies' wonderful result that these three notions coincide. In other words, a set has lowest possible information content iff it has no derandomization power iff it has no compression power. We examine several other properties of the K-trivial sets, including the fact that they are very close to being computable, and provide further characterizations of them, in terms of other notions of randomness-theoretic weakness and the important concept of a *cost function*.

In Chapter 12 we study lowness and triviality for other notions of randomness, such as Schnorr and computable randomness. For instance, we prove results of Terwijn and Zambella, and Kjos-Hanssen, Nies, and Stephan, characterizing lowness for Schnorr randomness in terms of traceability, and Nies' result that there are no noncomputable sets that are low for computable randomness. We also study the analog of K-triviality for Schnorr randomness, including a characterization of Schnorr triviality by Franklin and Stephan.

Chapter 13 deals with algorithmic dimension. Lutz realized that Hausdorff dimension can be characterized using martingales, and used that insight to define a notion of *effective Hausdorff dimension*. This notion turns out to be closely related to notions of partial randomness that allow us to say, for example, that certain sets are $\frac{1}{2}$-random, and also has a pleasing and useful characterization in terms of Kolmogorov complexity. We also study *effective packing dimension*, which can also be characterized using Kolmogorov complexity, and can be seen as a dual notion to effective Hausdorff dimension. Algorithmic dimension is a large area of research in its own right, but in this chapter we focus on the connections with computability theory. For instance, we can formulate a computability-theoretic version of the quite natural question of whether randomness can be extracted from a partially random source. We prove Miller's result that there is a set of positive effective Hausdorff dimension that does not compute any set of higher effective Hausdorff dimension, the result by Greenberg and Miller that there is a degree of effective Hausdorff dimension 1 that is minimal (and therefore cannot compute a 1-random set), and the contrasting result by Zimand that randomness extraction *is* possible from two sufficiently independent sources of positive effective Hausdorff dimension. We also study the relationship between building sets of high packing dimension and array computability, and study the concept of *Schnorr dimension*. In the last section of this chapter, we look at Lutz' definition of dimension for finite strings and its relationship to Kolmogorov complexity.

The final three chapters cover further results relating randomness, complexity, and computability.

One of the byproducts of the theory of K-triviality has been an increased interest in notions of lowness in computability theory. Chapter 14 discusses the class of *strongly jump traceable* sets, a proper subclass of the K-trivials with deep connections to randomness. In the computably enumerable case, we show that the strongly jump traceable c.e. sets form a proper subideal of the K-trivial c.e. sets, and can be characterized as those c.e. sets that are computable from every ω-c.e. (or every superlow) 1-random set. We also discuss the general (non-c.e.) case, showing that, in fact, *every* strongly jump traceable set is K-trivial.

In Chapter 15 we look at Ω as an operator on Cantor space. In general, our understanding of operators that take each set A to a set that is c.e. relative to A but does not necessarily compute A (as opposed to the more usual "computably enumerable in and above" operators of computability theory) is rather limited. There had been a hope at some point that the Omega operator might turn out to be degree invariant, hence providing a counterexample to a long-standing conjecture of Martin that (roughly speaking) the only degree invariant operators on the degrees are iterates of the jump. However, among other results, we show that there are sets A and B that are equal up to finite differences, but such that Ω^A and Ω^B are relatively 1-random (and hence Turing incomparable). We also establish several other properties of Omega operators due to Downey, Hirschfeldt, Miller, and Nies, including the fact that almost every real is Ω^A for some A, and prove Miller's results that a set is 2-random iff it has infinitely many initial segments of maximal prefix-free Kolmogorov complexity.

Chapter 16 is devoted to the relationship between Kolmogorov complexity and c.e. sets. We prove Kummer's theorem characterizing the Turing degrees that contain c.e. sets of highest possible Kolmogorov complexity, Solovay's theorem relating the complexity of describing a c.e. set to its enumeration probability, and several results on the complexity of c.e. sets naturally associated with notions of Kolmogorov complexity, such as the set of nonrandom strings (in various senses of randomness for strings).

As mentioned in the preface, we have included several open questions and unexplored research directions in the text, which are referenced in the index under the heading *open question*.

Unlike the ideal machines computability theorists consider, authors are limited in both time and space. We have had to leave out many interesting results and research directions (and, of course, we are sure there are several others of which we are simply unaware). There are also entire areas of research that would have fit in with the themes of this book but had to be omitted. One of these is the uses of algorithmic randomness in reverse mathematics. Another, related one, is the growing body of results on converting classical "almost everywhere" results in areas such as probability and dynamical systems into "for all sufficiently random" results (often precise ones, saying, for instance, that a certain statement holds for all 2-random but not all 1-random real numbers). Others come from varying one of three ingredients of algorithmic randomness: the spaces we consider, the measures on those spaces, and the level of effectiveness of our notions. There has been a growing body of research in extending algorithmic randomness to spaces other than Cantor space, defining for example notions of random continuous functions and random closed sets. Even in the context of Cantor space, we only briefly discuss randomness for measures other than the uniform (or Lebesgue) measure (although, for computable continuous measures at least, much of the theory remains unaltered). The interaction between randomness and complexity theory is a topic that could easily fill a book this size by itself, but there are parts of it that are particularly close to the material we cover. Randomness in the context of effective descriptive set theory has also begun to be investigated.

Nevertheless, we hope to give a useful and rich account of the ways computability theorists have found to calibrate randomness for individual elements of Cantor space, and how these relate to traditional measures of complexity, including both computability-theoretic measures of relative computational power such as Turing reducibility and notions from algorithmic information theory such as Kolmogorov complexity. Most of the material we cover is from the last few years, when we have witnessed an explosion of wonderful ideas in the area. This book is our account of what we see as some of the highlights. It naturally reflects our own views of what is important and attractive, but we hope there is enough here to make it useful to a wide range of readers.