The Book Review Column[1]

by Frederic Green

Department of Mathematics and Computer Science
Clark University
Worcester, MA 01610
email: fgreen@clarku.edu

In this column, we review these two books, each unique in its field:

1. **Complexity Dichotomies for Counting Problems – Volume 1: Boolean Domain**, by Jin-Yi Cai and Xi Chen. An expert and thorough introduction to this vibrant field of research. Review by Frederic Green.

2. **The Geometry of Uncertainty**, by Fabio Cuzzolin. An extensive treatise on belief and uncertainty theory. Review by Erick Galinkin.

Please contact me to write a review! Choose from among the books listed on the next page. Or choose one of your own. The latter is preferable (and quicker) in the current circumstances, as I can then ask the publisher to forward it directly to you.

BOOKS THAT NEED REVIEWERS FOR THE SIGACT NEWS COLUMN

Algorithms

1. *Algorithms and Data Structures Foundations and Probabilistic Methods for Design and Analysis*, by Helmut Knebl
2. *Algorithms and Data Structures*, by Helmut Knebl
3. *Beyond the Worst-Case Analysis of Algorithms*, by Tim Roughgarden

Computability, Complexity, Logic

1. *Applied Logic for Computer Scientists: Computational Deduction and Formal Proofs*, by Mauricio Ayala-Rincón and Flávio L.C. de Moura.
2. *Descriptive Complexity, Canonisation, and Definable Graph Structure Theory*, by Martin Grohe.
3. *Semigroups in Complete Lattices*, by P. Eklund, J. Gutiérrez García, U. Höhle, and J. Kortelainen.

Miscellaneous Computer Science

1. *Elements of Causal Inference: Foundations and Learning Algorithms*, by Jonas Peters, Dominik Janzing, and Bernhard Schölkopf.
2. *Partially Observed Markov Decision Processes,* by Vikram Krishnamurthy
3. *Statistical Modeling and Machine Learning for Molecular Biology*, by Alan Moses
4. *Language, Cognition, and Computational Models,* Theirry Poibeau and Aline Villavicencio, eds.
5. *Computational Bayesian Statistics, An Introduction,* by M. Antónia Amaral Turkman, Carlos Daniel Paulino, and Peter Müller.
6. *Variational Bayesian Learning Theory,* by Shinichi Nakajima, Kazuho Watanabe, and Masashi Sugiyama.
7. *Knowledge Engineering: Building Cognitive Assistants for Evidence-based Reasoning*, by Gheorghe Tecuci, Dorin Marcu, Mihai Boicu, and David A. Schum.
8. *Quantum Computing: An Applied Approach*, by Jack D. Hidary

Discrete Mathematics and Computing

1. *Mathematics in Computing*, by Gerard O'Regan
2. *Understand Mathematics, Understand Computing – Discrete Mathematics That All Computing Students Should Know*, by Arnold L. Rosenberg and Denis Trystram

Cryptography and Security

1. *Computer Security and the Internet: Tools and Jewels*, by Paul C. van Oorschot

Combinatorics and Graph Theory

1. *The Zeroth Book of Graph Theory: An Annotated Translation of Les Réseaux (ou Graphes) – André Sainte-Laguë (1926)*, translated by Martin Charles Golumbic
2. *Finite Geometry and Combinatorial Applications*, by Simeon Ball
3. *Combinatorics, Words and Symbolic Dynamics,* Edited by Valérie Berthé and Michel Rigo

Programming etc.

1. *Formal Methods: An Appetizer*, by Flemming Nielson and Hanne Riis Nielson
2. *Sequential and Parallel Algorithms and Data Structures*, by P. Sanders, K. Mehlhorn, M. Dietzfelbinger, R. Dementiev

Miscellaneous Mathematics

1. *Algebra and Geometry with Python*, by Sergei Kurgalin and Sergei Borzunov.

Review of[2]
Complexity Dichotomies for Counting Problems
Volume 1: Boolean Domain
by Jin-Yi Cai and Xi Chen
Cambridge University Press, 2017
461 pages, Hardcover, Price \$163

Review by
Frederic Green (fgreen@clarku.edu)
Department of Computer Science, Clark University

An Introduction:
This book is a definitive account of a field of research that has blossomed remarkably over the past two decades. Many of its chief developments are presented here, indeed in many cases having been obtained by its authors.

The broad context is this: It is an understatement to say it is widely believed that $\mathsf{P} \neq \mathsf{NP}$. It has also been known for a long time that *if* that is the case, then there are sets in NP that are neither in P nor NP-complete. One may also consider complexity classes of *counting problems*, typified by the class $\#\mathsf{P}$, for which an important complete problem is $\#\mathsf{SAT}$, the problem counting the number of satisfying assignments of a boolean formula. For such function classes a fact analogous to that regarding P vs. NP holds (indeed, via the same proof [Lad75]), namely that assuming (e.g.) $\mathsf{FP} \neq \#\mathsf{P}$, there are problems in $\#\mathsf{P}$ that are neither in FP nor $\#\mathsf{P}$-complete. However, for a large natural class of counting problems, every problem in the class is either $\#\mathsf{P}$-complete or solvable in polynomial time, a so-called "dichotomy." For boolean constraint satisfaction problems, dichotomies have long been known [Sha78] separating into either P or NP-complete. The study of dichotomies for counting problems is the driving theme of this book.

This particular line of research has an interesting history, some of it recounted here in the context of the book's goals and terminology. Back in 2002, Valiant published a seminal paper [Val02] exhibiting a large class of quantum circuits which can be simulated classically in polynomial time. Since the probability amplitudes associated with quantum circuits comprise a sum of exponentially many terms (each of which, in general, is a product of complex numbers) as a function of the input size, computing these amplitudes naturally leads to classical problems involving sums of products. The polynomial time simulation was achieved by representing the circuit as a graph, and reducing the computation of the amplitudes to counting the number of perfect matchings in that graph. Thus the gates of a quantum circuits are modeled by graphical gadgets called *matchgates* (more of which, see below). The reduction to counting the number of perfect matchings was achieved by relating it to the Pfaffian of a matrix (a quantity whose square is the determinant for an even-dimensional matrix, and 0 for odd). Using the fact that the Pfaffian is computable in polynomial time, it followed that the probability amplitudes were poly-time computable.

The reduction used in [Val02] was a standard Turing reduction. In a subsequent paper [Val08], Valiant introduced the novel idea of *holographic reductions*, which are used to derive what he called *holographic algorithms*. The "holographic" designation is not due to any analog to the encoding of the whole image from part of its interference pattern as in a hologram, but rather from the fact that "destructive interference," manifested as *cancellation*, plays an instrumental role. While one may view holographic reductions as deriving from the fact that a quantum circuit entails a unitary transformation and therefore leads to useful cancellations, they can also be understood (and vastly generalized) quite apart from quantum computing.

To see (roughly) how holographic reductions work, consider the following setting. We start with a k-regular weighted graph $G = (V, E)$, along with an assignment of a function f_v of arity k to each vertex v. In general, one might set up a mapping between a *set* of functions and the set of vertices, but for the sake of simplicity (and a particular example), suppose it is just one function. The inputs of that function at each v are given by the weights of the incident edges. To further narrow down to an example, suppose the edge weights are all 1 or 0 and that our function is $f_v = \text{EXACT-ONE}_k$, which returns 1 if exactly one of the k inputs is 1, and 0 otherwise[3]. That is, it will output 1 if exactly one of the incident edges is assigned weight 1. Now take the product over all vertices of the value of that function. It is not hard to see that the product will equal 1 iff the weight 1 edges specify a perfect matching in the graph. Now sum that product over all assignments of weights to edges. The result is the number of perfect matchings. Symbolically, we can represent this sum as follows:

$$\sum_{\sigma: E \to \{0,1\}} \prod_{v \in V} f_v(\sigma|_{E(v)}),$$

where σ is the function that assigns weights to edges, and $f_v(\sigma|_{E(v)})$ denotes the function evaluated at vertex v with the weights assigned according to σ. We have just described an example of what is called a *signature grid* (the graph G along with the assignments of functions to vertices), and the above sum of products is called the *Holant* of that signature grid. Thus the Holant can express important quantities like the number of perfect matchings; it is of course much more general and can serve to express numerous other quantities, e.g., the number of cycle covers. Now, without yet saying exactly what holographic reductions are, their importance rests on the fact that *the Holant is invariant under holographic reductions.*

One can represent the functions at each vertex by a vector, called here a *signature* (the source of the term "signature grid" above), where the components of the vector correspond to the values of the function, indexed by its arguments. For example, the function EXACT-ONE_2 would be represented by the vector $(0, 1, 1, 0)$, or EXACT-ONE_3 by $(0, 1, 1, 0, 1, 0, 0, 0)$. Avoiding excessive technical detail (including a non-trivial proof), holographic transformations are matrices acting on these vectors in such a way as to keep the Holant invariant. In the case of bipartite graphs, the only requirement of these matrices is that they be invertible. As we allow these matrices to have complex entries, this includes unitary matrices, one relic of the quantum origins of this technique. For general graphs, they must be orthogonal (but still over $\mathbb{C}$). And note that these matrices are general enough to allow the cancellation promised earlier.

This brings us to matchgates, which play a central role in holographic algorithms. A matchgate is a graph (understood to be used as a gadget in a reduction) with a number of "external" vertices, which are assigned a standard order. We associate with the matchgate a signature, where the indexing is based on the ordering of the external vertices. This signature is based on a quantity similar to the Holant, but now summing products of the weights over perfect matchings of the graph. Recalling that holographic reductions preserve the Holant, we can transform a signature grid to another one which replaces the original vertices with matchgates, *if* those matchgates can "realize" the signatures associated with the vertices of the original grid. We can't naively substitute the original vertices with matchgates with appropriate numbers of external vertices, because certain signatures cannot be realized by matchgates. However, by choosing judicious holographic transformations, in many instances we *can* transform the signatures to ones that *are* realizable by matchgates, in which case the substitution works. Finally, given a polynomial-time algorithm for solving the problem on the transformed signature grid, we obtain an algorithm for the original problem, hence "holographic algorithm."

[3]There being only one function to choose from, the subscript in f_v is, in this case, superfluous.

The algorithm that fits the bill for counting planar perfect matchings is variously known as Kasteleyn's Algorithm or the FKT Algorithm, after Kasteleyn, and Fisher and Temperley, the physicists who independently discovered it in 1961 [Kas61, TF61]. It was at that time used to understand the statistical mechanics of "dimer models" in statistical mechanics. This comprises the "fully frustrated" case of the Ising model, which, it turns out, is equivalent to counting the number of domino tilings of a region of the plane, and also to counting the number of perfect matchings in a planar graph. While it had been known for a long time, the widespread applicability of the algorithm to such a broad collection of counting problems is a more recent development. In fact, it is in a certain sense "universal" for planar #CSP problems (defined below), a theorem which is one of the highlights of this book. (While this universality does not hold for the Holant, dichotomies can nevertheless be proved in that setting, as is done in Chapter 9 of the book. A little more detail on this below.)

Holographic algorithms present us with some surprises that, at least on the surface, seem somewhat bizarre. A notable example is the following:

$\#_7$PL-RTW-MON-3CNF
INPUT: A planar 3CNF Boolean formula where each variable appears positively and in exactly two clauses (planar, read-twice, monotone, 3CNF).
OUTPUT: The number of satisfying assignments modulo 7.

This is solvable in polynomial time. Perhaps not surprisingly, dropping the mod 7, it is $\#\mathsf{P}$-complete. But far more curiously, simply replacing the 7 with a 2 (for example), we obtain $\#_2$PL-RTW-MON-3CNF, which is $\oplus\mathsf{P}$-complete! In fact, $k = 7$ is the *only* value of k which guarantees that $\#_k$PL-RTW-MON-3CNF is in P via the matchgate technique.

So what is going on here? It turns out that, for a given problem, finding a set of matchgates that realize given signatures (and, hence, determining whether or not a holographic algorithm exists) reduces to solving a set of algebraic equations. They are not easy to solve. Valiant referred to them as "an apparently onerous set of constraints," resulting in an enumeration that can lead to "accidental or freak objects." He even called the resulting algorithms *accidental* [Val06]. Of course, as he knew then, they only appear accidental. To paraphrase Einstein, mathematics is subtle but not malicious. Fortunately, this puzzle is addressed in this book. This is one spoiler I will leave out of my review. You'll have to read it to find out!

The Contents:
The foregoing was both a brief historical account and a summary of a substantial portion of the book. Before proceeding, however, recall that the subtitle of this book is "Volume I: Boolean Domain," i.e. the variables only take on boolean values (whereas constraint functions can take on complex values). The upcoming Volume II will address more general domains in which variables can take on any values in a finite domain. Throughout this review, I have been (and will continue to be) cavalier about these distinctions.

That being said, here is more about the presentation: The first chapter quickly gets into the definitions of the various problems, i.e., spin systems, constraint satisfaction problems ("CSPs," the counting version being denoted by #CSP) and Holants. We also encounter holographic transformations and the fact that they leave quantities such as the Holant invariant. Fibonacci gates, an alternate primitive to matchgates which can also be used to obtain holographic algorithms, are introduced in Chapter 2 in order to prove a dichotomy theorem for "Holant* problems," in which unary functions are allowed for free, and in the case where all signatures are symmetric functions. This and most of the remaining chapters lay out a theory which leads to greatly generalized results later on, while proving various dichotomies along the way. Thus Chapter 3 proves a dichotomy for #CSPs, importantly allowing for asymmetric constraints; most of the results up to chapter 9 are for *symmetric* constraints. The theory of matchgates is the focus of Chapter 4. Here we see that matchgates realize signatures that obey a remarkable set of identities ("matchgate identities" or "MGIs"). It turns out that the MGIs are equivalent to the realizability of signatures by matchgates, a vital tool in later proofs. This chapter also solves the mystery of $\#_7$PL-RTW-MON-3CNF. Moving on, Chapter 5 proves dichotomies for 2-spin systems on k-regular graphs, which may be viewed as restrictive cases of #CSPs, or alternatively bipartite Holant problems. Both Chapters 3 and 5 use results about Holant* problems to obtain dichotomies. Chapter 6 further explores the interaction between Holant problems, #CSPs, and spin systems by showing how one can go in the other direction and obtain, *from* #CSPs and 2-spin systems, dichotomies for "Holant^c" problems, a class intermediate between Holant* and Holant problems. Chapter 7 proves dichotomies for Holant problems with symmetric constraint functions, with no assumptions about allowing signatures for free. Thus the results here are built on and also imply certain key results of Chapters 2, 3 (but then restricted to symmetric constraints), and 6. Chapter 8 proves a dichotomy for the #CSPs of Chapter 3 in which the underlying graph is planar and the constraints symmetric, which gives an extended criterion for determining if the problem is in P on planar structures. Chapter 9 marks the culmination of the first 9 chapters, proving a dichotomy for planar Holant problems with symmetric constraints, with no assumptions about freely available signatures. Another surprise comes up here: Results of previous chapters hint that all these problems would fall into one of three categories, namely (1) P-time solvable for general graphs, (2) P-time solvable over planar graphs but $\#\mathsf{P}$-hard over general graphs, or (3) $\#\mathsf{P}$-hard over planar graphs. Furthermore, those results suggest that those in category (2) are a consequence of holographic reductions to Kasteleyn's algorithm. However, while the categorization remains true for planar Holant problems, the P-time algorithms do not result only from Kasteleyn's. Finally, Chapter 10 addresses *asymmetric* constraints, for which not many results are known. Generalizations of results from Chapters 2 and 8 are presented, although in the latter case an outline is given of the full proof (which can be found in the literature, see the sources [CFG22, CF22]). Interestingly, the results are obtained by reducing to the symmetric case.

My Opinion:

When I first glanced at the book, I had some trepidations that it might be a bit too specialized for a general (TCS) audience. After getting into it, I quickly realized how wrong I was. The challenges it presents are eminently worth overcoming, and the authors have provided all the means to overcome them. There's no doubt in my mind that anyone interested in computational complexity would benefit greatly by reading it. More general readers in computer science will also find much that is fascinating, surprising, and even inspiring, to say nothing of useful.

The logical organization from the first chapter to the last, progressing from more special to increasingly general results, interspersed with a rich set of exercises, exhibits great care on the part of the authors to make this difficult topic understood. Given these pedagogical virtues, professors should seriously consider courses based on the book. It would be perfect for one semester at the graduate level. The prerequisites are "modest." That is in quotes because, while no knowledge of computational complexity or quantum computing is assumed, and it seems that decent courses in undergraduate algebra "done right" would be enough, a healthy dose of mathematical maturity is mandatory. Given that, it is accessible to determined undergraduates.

To sum up, this book is a foundational and authoritative introduction, the first and so far the only unified treatment of these topics bound between two covers. It is a remarkable and indispensable book about a deep and important topic. I cannot recommend it highly enough.

And we can look forward to Volume II, which should include such sweeping results as the #CSP dichotomy on general domains, for which the authors have just won the Gödel Prize in 2021!

References:

[CFG22] J.-Y. Cai, Z. Fu, H. Guo, T. Williams, "FKT is not universal – A planar holant dichotomy for symmetric constraints," in *Theory Comput. Syst.* **66**, pp. 143–308 (2022).

[CF22] J.-Y. Cai and Z. Fu, "Holographic algorithm with matchgates is universal for planar #CSP over boolean domain," in *SIAM Journal on Computing* **51**(2), pp. STOC17-50–STOC17-151 (2022).

[Kas61] P. W. Kasteleyn, "The statistics of dimers on a lattice," in *Physica* **27a**: pp. 1209–1225 (1961).

[Lad75] R. Ladner, "On the structure of polynomial time reducibility," in *Journal of the ACM* **22**(1), pp. 155–171.

[Sha78] T. K. Schaefer, "The Complexity of Satisfiability Problems," in *10th Annual ACM Symposium on Theory of Computing*, pp. 216226 (1978).

[TF61] H. N. V. Temperley and M. E. Fischer, "Dimer problems in statistical mechanics–An exact result," in *Philosophical Magazine* **6**: pp. 1061–1063 (1961).

[Val02] L. Valiant, "Quantum circuits that can be simulated classically in polynomial time," in *SIAM Journal on Computing* **31**(4): pp. 1229–1254 (2002).

[Val06] L. Valiant, "Accidental Algorithms," in *Proceedings of the IEEE 54th Annual Symposium on Foundations of Computer Science*, pp. 509-517 (2006).

[Val08] L. Valiant, "Holographic algorithms," in *SIAM Journal on Computing* **37**(5): pp. 1565–1594 (2008).

Review of[4]
The Geometry of Uncertainty
by Fabio Cuzzolin
Springer-Verlag, 2021
850 Pages, Hardcover, $279.99

Review by
Erick Galinkin (erick_galinkin@rapid7.com)
Rapid7

1 Introduction

The Geometry of Uncertainty is unlike any book on mathematics and computer science I've ever read. It's certainly not a textbook in the traditional sense – there are no exercises, and very little is presented as the "right way" to do something. In many ways, *The Geometry of Uncertainty* is like a survey paper: it critically analyzes decades of research and while some of the author's preferences are reflected, many perspectives are presented with little direct guidance about which approach ought to be favored. People have long used the phrase "The Bible of X" to describe the most well-known or well-respected text in a field. However, calling *The Geometry of Uncertainty* "The Bible of Evidence Theory" is truly apt in the sense that you come into the book with questions, and you leave the book with parables, historical insight, and a whole host of new, better informed questions. Truly, the book is not merely an introduction to belief function theory, it is instead a truly comprehensive view of an entire field of study from its conception up to the most bleeding edge research.

2 Summary of Contents

The book is broken into four overarching sections followed by a brief fifth section identifying productive avenues for future work. Finally, the book is capped off by, per a Springer representative, "one of the largest bibliographies we have ever published".

After a relatively short introduction to probability theory as it exists today and some of the shortcomings therein, we move into the first section on the various theories of uncertainty. This first section provides an introduction to uncertainty and provides some tools and techniques for dealing with uncertainty and evidence. This first section is comprised of six chapters which begin by introducing belief functions – a function that returns a set rather than a single value – and moves through the subsequent chapters by discussing the conflicts, conclusions, history, and approaches of various researchers who have advanced belief function theory. Some of the names in these chapters – Dempster, Shafer, Pearl – are quite familiar to anyone in this book's target audience, but quite a few names that are likely less familiar to non-experts are also discussed at length. This first section of the book leaves readers with a toolbox full of combination and conditioning methods, a variety of entropy and conflict measures, and a firm idea of the strengths and weaknesses of each. However, it is left to the reader to decide which of these tools are valuable in what situation and what their preferred methods are.

The second section builds on this initial introduction by introducing geometry into our belief functions. This section reformulates much of the first section in geometric language. Critically, much of the language here uses concepts that are difficult to introduce within the confines of a book that does not deal with geometry as its primary subject matter and so it would behoove readers to familiarize themselves with simplices, manifolds, fiber bundles, and a bit of differential geometry before delving into this section. The eighth chapter of the book, in particular, compellingly describes the geometry of Dempster's rule. In this reviewer's opinion, the geometric description provides excellent intuition for the applications of Dempster's rule.

The third section advances our geometric conception of belief functions and probability. In particular, the section begins with two families of probability transforms and a discussion of not only how and where to use them – but importantly, why we want them at all. This section continues on to provide us approximations for mass and belief spaces.

The fourth section consists of only two chapters, but carries some of the weightiest material in the book. The first chapter of section four deals with geometric condition under various L_p norms using our geometric framework. The second chapter, chapter 16, harkens back to the epistemic transforms introduced in chapter 12 and provides tools for decision making using these epistemic transforms and the geometry introduced throughout sections 2 and 3.

The final section of the book consists of a 66 page chapter that reads like a world-class research agenda. Although the particular topics are likely biased by Cuzzolin's own research interests, they remain extremely well motivated and Cuzzolin provides readers no shortage of references on which to build their research program. As was mentioned earlier, this book provides sufficient amounts of background and references for dozens of PhDs, grants, or research groups. Indeed, the book offers 2137 references to papers and books throughout and each of these citations is there for the taking should someone wish to turn to the source material following Cuzzolin's presentation.

3 Chapter Highlights

Chapter 1 provides introduction and motivation for the rest of the book, describing the mathematical notions of probability by way of Kolmogorov's axioms, and explaining the competing interpretations of probability – the well-trodden Frequentist vs. Bayesian argument. Cuzzolin then sets up a truly salient issue, and one that often gets left out of other texts on probability: there are significant deficiencies in both interpretations of probability. Moreover, there are cases where Frequentist and Bayesian inference will reach opposite conclusions given observed data. The chapter concludes with a discussion of uncertainty and the open debate on how to best model uncertainty – in Cuzzolin's view, via the use of evidence and belief functions.

Chapters 2 and 3 set up the necessary framing for Dempster's theory of evidence using a motivating example of a murder trial where different evidence is introduced. In the interest of brevity, I will paraphrase the example here:

Imagine an investigation into a murder. Let our set of suspects be denoted $\Theta = \{\text{Peter}, \text{John}, \text{Mary}\}$ where Peter and John are men, and Mary is a woman. Consider testimony Ω from a witness asserting that they saw a man at the scene, but they had been drinking, so they can only be 80% sure. This motivating example provides us with a probability assignment $P(\text{Saw a man}) = 0.8, P(\text{Didn't see a man}) = 0.2$. How do we assign this probability to Θ?

Using the theory of belief functions, we can create a multivalued mapping $\Gamma : \Omega \to 2^\Theta$ that induces a mass assignment $m : 2^\Theta \to [0, 1]$ on the power set of Θ. This allows us to map $\Gamma(\text{Saw a man})$ onto $\{\text{Peter}, \text{John}\}$ and assign the remaining probability mass to Θ, yielding this basic probability assignment $m(\{\text{Peter}, \text{John}\}) = 0.8)$ and $m(\Theta) = 0.2$. This is important because in traditional probability theory, we would have to choose a distribution over $\{\text{Peter}, \text{John}\}$, and either assign the remaining probability mass to

Mary or allocate it uniformly over the 3 suspects. The value of this is crucial – we do not, at this point, have any real reason to believe that Mary should be a suspect, but our suspicion of John and Peter is bounded by the probability that the witness actually saw a man at the scene. This leads us to the definition of belief and plausibility, where for some subset $A \subseteq \Theta$, we define the belief Bel and plausibility Pl by:

$$Bel(A) = \sum_{B \subseteq A} m(B)$$

$$Pl(A) = 1 - Bel(\overline{A})$$

where $\overline{A}$ is the complement of A in Θ. These chapters also introduce Dempster's rule of combination, which allow us to derive the aforementioned Γ inducing basic probability assignments on Θ when we have multiple sources of evidence, that is, multiple compatible Ω_i that we wish to map onto 2^{Θ}, even when these evidence conflict with one another.

Chapters 4 and 5 provide the remaining tools necessary to reason with belief functions and use them in a computational context – for ranking, clustering, classification, and so on. These chapters deal with measuring uncertainty in different ways, how to efficiently compute belief functions, and how we can make inferences and decisions under uncertainty. One interesting highlight that is left out of many treatments is Cuzzolin directly confronting the critiques of Dempster's rule of combination. In lieu of offering Dempster's rule or some alternative as "the one true rule," he instead provides insight into why there are a variety of combination rules, and how they relate to one another. In this reviewer's opinion, the overview of critiques, modifications, and suggestions is remarkably useful for researchers seeking to use the combination rules in practice.

Chapter 6 concludes the first part of the book, and covers capacities, fuzzy sets, rough sets, higher-order probabilities, and logic. This chapter demonstrates similarities, equivalences, and interpretations of the relationship between belief functions and these other objects, concluding with a variety of other uncertainty theories and formalisms. At 320 pages, these 6 chapters alone would likely have provided an excellent introduction to the theory and history of belief functions.

Chapters 7 and 8 delve into the geometry of belief functions and Dempster's rule. An important caveat about these chapters: Although Cuzzolin does his best to cover the geometric background necessary to understand the material, he is constrained to a finite space and so it is more of a refresher than a first course. Since probability distributions on finite domains can be represented as points on a simplex of n vertices, where n is the size of the sample space, we can leverage the language and machinery of convex geometry on belief functions. These chapters establish a framework for the second part of the book, motivated by this representation, where instead of examining only the power set of Θ, we consider vector representations in $\mathbb{R}^{|2^{\Theta}|}$, yielding a vector space where each vector is potentially a belief function. Returning to the Möbius transform of chapter 2, we are shown that determining whether a vector in the space corresponds to a belief function is exactly checking whether the transform of that vector satisfies the axioms of basic probability assignments. This reframing of belief functions and probability assignments is given some rigorous tools to limit the entire space of 2^{Θ}-dimensional vectors to merely the subset that constitutes the belief space, and a geometric rule of combination is derived.

Chapters 9 and 10 work to reframe earlier concepts on plausibility, commonality, and possibility into a geometric framework. Chapter 9 develops the equivalence of plausibility and commonality functions from chapter 2 in more detail and adds a third equivalent model in the form of a geometric framework. Chapter 10 deals with possibility and the consonant subspace, allowing for a geometric development of the theory of possibility. These chapters conclude with a handful of open research questions interested in extending certain concepts to notions of consistency and fuzzy sets.

Chapters 11 and 12 cover two families of probability transforms, the affine and epistemic, respectively. These chapters offer a geometric retelling of the relationship between belief functions and probability with respect to evidence covered in chapter 4, and develop approximation and conditioning approaches used in the following chapters.

Chapters 13 and 14 deal with problems of approximation, demonstrating methods for computing consonant and consistent approximations under L_1, L_2, and L_∞ distances in the belief space.

Chapter 15 deals with geometric conditioning, and some of the technical machinery likely requires readers to consult another reference, though careful readings of the proofs may be sufficient for some readers with sufficient mathematical maturity. The outline of future research at the end of this chapter raises some questions that seem particularly fruitful for future study. The question on whether Dempster conditioning is a special case of geometric conditioning seems especially useful for understanding when to use particular conditioning rules.

Chapter 16 provides tools for decision making with epistemic transforms, bringing together the tools introduced in chapters 7 through 15. This chapter shows how the simplicial form and probability transformations in the geometric framework can be used to derive a generalization of the Transferable Belief Model, and useful applications within game theory.

Chapter 17, the final chapter of the book sets forth open questions, directions for future work, and provides the final bit of motivation for this book.

4 Opinion

One firm critique that must be leveled at the book is the lack of an index. Given the size and scope of the book, one can imagine wanting to look up a particular reference or definition and the lack of an index means sifting through possibly hundreds of pages to find the relevant information – perhaps something that is easier to deal with in electronic versions of the text. For this reviewer's part, I found that noting down in the back of a notebook what pages certain information I'd wish to look up again made the book much easier to use, especially when I found references to information from several chapters earlier that I'd forgotten.

For those of us looking for something more general than traditional measure-theoretic probability, the material in the book represents an excellent opportunity to move beyond the bounds of Bayesian analysis when something more general may be needed. Certainly, there are accolades and advances that can and will be earned exploring the questions and contradictions presented in the book – many of which are presented as research questions at the end of chapters. However, those who would most benefit from the book are researchers, those looking for a PhD topic, and the truly interested and dedicated. The book has some shortcomings – some concepts, like those in Chapters 7, 8, and 15 are a bit technical to grasp without a background in Topology and Geometry; and there are more accessible materials on the non-geometric theory of belief functions, including lecture slides online made available by Cuzzolin himself. That said, I do not know of any textbook that covers the geometric aspects Cuzzolin details in the latter half of the book, making it well worth the journey for interested researchers. For those who want to truly understand the theory of belief functions from Dempster's first publication to the current research and open questions, you can do no better than The Gospel of Cuzzolin. This again sets Cuzzolin's work in rarified company, as the final chapter of the book does not merely set forth an agenda for future applications, or even potential further study in the same field. Rather, the final chapter of the book seeks to inspire the reader and let them know that there is more research to come – and that you're welcome to be a part of it.