Review of [1]

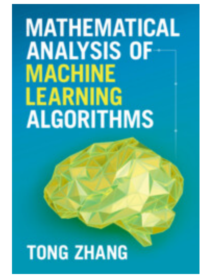

Mathematical Analysis of Machine Learning Algorithms
by **Tong Zhang**

Cambridge University Press, 2023
Hardback, 479 pages, $54.99

Review by **Chinmay Hegde**
New York University

1 Overview

Progress in machine learning as an academic discipline has been relentless. Over the last two-plus decades, the field has evolved from being the purview of a small set of experts in academia and top industrial labs to impacting (arguably) all of computer science, data science, applied mathematics, and beyond. The center-of-gravity of scholarly research has been shifting from mathematical ideas, concepts, and theories towards concrete applications, systems, and real-world use cases.

Within this context, the new textbook *Mathematical Analysis of Machine Learning Algorithms* by Professor Tong Zhang is a *tour de force*. The book provides an introduction to a variety of mathematical tools that underpin modern machine learning techniques. It serves as a reminder that the foundations of machine learning have been (and continue to be) derived from mathematical analysis, and that we as a community should preserve and enrich these elements. In many ways, this book echoes earlier classical texts such as [HTFF09].

This advanced textbook, geared towards graduate students and doctoral candidates, encapsulates several broad families of concepts in machine learning theory. Zhang is a renowned machine researcher, and I have been reading his papers since my early graduate student years. The breadth and depth of coverage mirror Zhang's own research contributions, reflecting an illustrious career spent at the forefront of machine learning theory. This also adds an extra layer of authority to the text, as readers may benefit from insights gleaned from several decades of research.

The book's scope is ambitious. It surveys a number of topics, each representing a distinct neighborhood in the landscape of machine learning theory. Starting with Probably Approximately Correct (PAC) learning and the statistical elements of generalization theory, the textbook moves onto kernel methods and nonlinear data analysis. Generalized linear models, which form the backbone of many predictive algorithms used in practice, are examined next. The treatment of neural networks is brief but offers a peek into the dizzying array of recent theoretical developments that have attempted to probe deeper into the building blocks of modern artificial intelligence. The last quarter of the book ventures into the fields of online learning and reinforcement learning, both areas being at the forefront of several current research agendas.

What struck me about this book is its unwavering focus on mathematical rigor. Zhang makes no concessions to the reader and does not fall prey to handwavy arguments and simplification at the expense of mathematical precision. This approach is undoubtedly demanding, but I believe this will help readers gain a deeper understanding of the subject matter.

Early in the book, the author makes it clear that he will avoid lengthy descriptions of standard machine learning algorithms and applications. The reader is assumed to be familiar with the basics and does not need to be convinced why they should be interested in the techniques that are being discussed. This is a bold choice, but one that I like: I appreciated the author clearly flagging the book's intended role, which is that it is a complement to (rather than a replacement for) more application-oriented textbooks.

For the same reason, the book's approach may present challenges for some readers. The dense mathematical content and theoretical focus requires a high level of mathematical maturity, perseverance, and discourse. This book is undoubtedly *not* meant for casual perusal, but rather a work that will reward thinking and reflection.

Despite this challenging style (or perhaps because of it!), *Mathematical Analysis of Machine Learning Algorithms* will greatly benefit those seeking to understand the theoretical foundations of machine learning at the highest level. It fills a distinct niche in the existing literature. I think the book will stand as a testament to the depth and maturity of machine learning as a field. It challenges readers to elevate their understanding of learning algorithms beyond mere applications to grasp the fundamental principles that may drive further innovation. For those willing to embark on this intellectual journey, the rewards will be substantial.

2 Book Contents

The book is organized into 18 chapters. I will take the liberty to loosely cluster these chapters into seven sections. Here is a brief overview of each section:

1. **Introduction and Probability Theory** (Chapters 1-2): The first couple of chapters lay the groundwork for the rest of the book. The author introduces the fundamental concepts of machine learning and reviews essential probability theory. These chapters ensure that readers will have the necessary mathematical foundation to absorb the more advanced topics that follow later in the book.

2. **Generalization Theory** (Chapters 3-6) The next several chapters focus on the theoretical question of why and how machine learning methods perform well when presented with new data. To address this, the book introduces the concepts of risk minimization, uniform convergence, hypothesis classes and their VC-dimension, covering numbers, and Rademacher complexity. These concepts are essential for any researcher interested in understanding the statistical limits of model generalization from training data to unseen examples.

3. **Stability and Model Selection** (Chapters 7-8) These chapters provide an alternative lens to view generalization, focusing algorithmic stability and model selection techniques. While classical generalization theory focuses on establishing risk bounds based on the complexity of the hypothesis class, algorithmic stability obtains risk bounds based on the learning algorithm. Model selection is essential for choosing the best model among various alternatives.

4. **Advanced Learning Techniques** (Chapters 9-11) The book then moves onto specific families of machine learning algorithms. The next few chapters cover theoretical elements of several powerful machine learning techniques: kernel methods, additive models, random feature models, and neural networks. It provides a theoretical treatment of these models, explaining their mathematical foundations and properties. Here is also where the author (briefly) touches on capacity questions of machine learning models, the Universal Approximation Theorem, and the intriguing phenomenon of double descent.

5. **Lower Bounds** (Chapter 12): Much of the book focuses on upper bounds on how well machine learning algorithms may perform; this chapter deviates from this narrative and asks the question: "What's the best possible performance achievable?", introducing information-theoretic tools such as Fano's Inequality. Understanding these lower bounds is crucial for recognizing the inherent limitations of learning and the hardness of certain problems.

6. **Sequential Learning and Online Methods** (Chapters 13-15) These chapters make the transition standard (offline) learning to sequential (online) learning settings. There is a brief interlude again where the author revisits the basics of sequential random variables, online learning algorithms, and posterior averaging techniques. All these are important for analyzing streaming data and adapting models in real-time.

7. **Decision Making Under Uncertainty** (Chapters 16-18) The final section of the book explores topics at the intersection of machine learning and decision theory. It covers multi-armed bandits, contextual bandits, and reinforcement learning (RL). This serves as a very accessible theoretical treatment of bandit and RL algorithms, which are vital for making decisions in uncertain environments while learning from feedback.

3 Critiques

This book is a remarkable achievement. However, as with any book with an expansive scope, there are certain aspects that are uniquely excellent and other aspects that could potentially be improved in future editions.

Let me start by highlighting the book's best aspects. I enjoyed the book's comprehensive coverage of the topic; it offers both a broad and deep exploration of machine learning theory, covering foundational concepts and advanced research ideas alike. I particularly enjoyed the last three chapters on online algorithms and reinforcement learning; apart from a small handful of great books such as [AJKS19], this area has received relatively lesser attention from the machine learning theory community and may benefit from more comprehensive treatises.

I already mentioned above that Zhang's commitment to mathematical precision and theoretical depth is commendable and will be immensely rewarding to the dedicated reader. The book also benefits from Zhang's extensive research contributions: it is filled with insights that bridge the gap between theoretical concepts and cutting-edge research.

While the book's strengths are numerous, I believe there are two main aspects where future editions could potentially be enhanced:

1. **Optimization Aspects**: This might be my own personal bias seeping through, but I felt that the book could benefit from a more comprehensive treatment of *optimization aspects* in

the context of machine learning. This is itself a vast area of theoretical study, but somewhat under-emphasized (in my view) within the contents of this book.

Research over the last ten years (highlighted in [NTS15], but likely dating back much earlier) has pointed to intricate relationship between optimization and generalization in machine learning, particularly through the lens of implicit bias. Expanding on this relationship could provide readers with the key understanding of how optimization strategies impact learning outcomes and generalization performance.

Addressing this subject as thoroughly as the other topics of this book will be no small feat. But I believe that this set of topics could further enrich the book's already comprehensive theoretical treatment and align it with current research trends.

2. **Chapter Organization and Narrative Flow**: The current sequence of chapters, while logical, may sometimes feel disconnected to readers. It took me a little bit of time to appreciate why Chapters 7-8 (Stability and Model Selection) and Chapter 12 were placed in their respective locations in the book.

   To improve this aspect, the author could consider introducing a set of recurring, real-world inspired problems or examples that thread through the different sections. These "running examples" could serve multiple purposes:

   - They would provide practical context for the theoretical concepts, helping readers understand the real-world implications of the mathematical principles discussed.
   - They could act as a unifying element, tying together different chapters and sections.
   - Such examples could make the material more engaging and accessible, especially for readers who are application-oriented.

4 Conclusion

The book stands as a monumental achievement, and Zhang deserves high praise for this work. It spans both advanced research and academic instruction, while also pushing the boundaries of our understanding of the mathematical foundations underlying modern machine learning techniques. Zhang's expertise shines through every page. I hope that it will solidify the foundational thinking of the next wave of machine learning researchers who will undoubtedly shape the future of the field.

No doubt there will be room for minor enhancements in future editions. But the current work is already an indispensable resource for graduate students, researchers, and anyone seeking a rigorous understanding of machine learning. This book is a testament to both the author's extensive knowledge and to the maturation of machine learning as a scientific and mathematical discipline. I will recommend it very highly to my colleagues and students.

References

[AJKS19] Alekh Agarwal, Nan Jiang, Sham M Kakade, and Wen Sun. Reinforcement learning: Theory and algorithms. *CS Dept., UW Seattle, Seattle, WA, USA, Tech. Rep*, 32:96, 2019.

[HTFF09] Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, volume 2. Springer, 2009.

[NTS15] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2015.