
Foreword

383 Machine learning is the latest in a long line of attempts to capture human
384 knowledge and reasoning into a form that is suitable for constructing ma-
385 chines and engineering automated systems. As machine learning becomes
386 more ubiquitous and its software packages become easier to use it is nat-
387 ural and desirable that the low-level technical details are abstracted away
388 and hidden from the practitioner. However, this brings with it the danger
389 that a practitioner becomes unaware of the design decisions and, hence,
390 the limits of machine learning algorithms. The enthusiastic practitioner
391 who is interested to learn more about the magic behind successful ma-
392 chine learning algorithms currently faces a daunting set of pre-requisite
393 knowledge:

- 394 • Programming languages and data analysis tools
- 395 • Large-scale computation and the associated frameworks
- 396 • Mathematics and statistics and how machine learning builds on it

397 At universities, introductory courses on machine learning tend to spend
398 early parts of the course covering some of these pre-requisites. For histori-
399 cal reasons, courses in machine learning tend to be taught in the computer
400 science department, where students are often trained in the first two ar-
401 eas of knowledge, but not so much in mathematics and statistics. Current
402 machine learning textbooks try to squeeze in one or two chapters of back-
403 ground mathematics, either at the beginning of the book or as appendices.
404 This book brings the mathematical foundations of basic machine learning
405 concepts to the fore and collects the information in a single place.

406 *Why Another Book on Machine Learning?*

407 Machine learning builds upon the language of mathematics to express
408 concepts that seem intuitively obvious but which are surprisingly difficult
409 to formalize. Once properly formalized we can then use the tools of math-
410 ematics to derive the consequences of our design choices. This allows us
411 to gain insights into the task we are solving and also the nature of intel-
412 ligence. One common complaint of students of mathematics around the
413 globe is that the topics covered seem to have little relevance to practi-
414 cal problems. We believe that machine learning is an obvious and direct
415 motivation for people to learn mathematics.

416 This book is intended to be a guidebook to the vast mathematical lit-
 417 erature that forms the foundations of modern machine learning. We moti-
 418 vate the need for mathematical concepts by directly pointing out their
 419 usefulness in the context of fundamental machine learning problems. In
 420 the interest of keeping the book short, many details and more advanced
 421 concepts have been left out. Equipped with the basic concepts presented
 422 here, and how they fit into the larger context of machine learning, the
 423 reader can find numerous resources for further study, which we provide at
 424 the end of the respective chapters. For readers with a mathematical back-
 425 ground, this book provides a brief but precisely stated glimpse of machine
 426 learning. In contrast to other books that focus on methods and models of
 427 machine learning (Mackay, 2003; Bishop, 2006; Alpaydin, 2010; Rogers
 428 and Girolami, 2016; Murphy, 2012; Barber, 2012; Shalev-Shwartz and
 429 Ben-David, 2014) or programmatic aspects of machine learning (Müller
 430 and Guido, 2016; Raschka and Mirjalili, 2017; Chollet and Allaire, 2018)
 431 we provide only four representative examples of machine learning algo-
 432 rithms. Instead we focus on the mathematical concepts behind the models
 433 themselves, with the intent of illuminating their abstract beauty. We hope
 434 that all readers will be able to gain a deeper understanding of the basic
 435 questions in machine learning and connect practical questions arising
 436 from the use of machine learning with fundamental choices in the mathe-
 437 matical model.

438 *Who is the Target Audience?*

439 As applications of machine learning become widespread in society we be-
 440 lieve that everybody should have some understanding of its underlying
 441 principles. This book is written in an academic mathematical style, which
 442 enables us to be precise about the concepts behind machine learning. We
 443 encourage readers unfamiliar with this seemingly terse style to persevere
 444 and to keep the goals of each topic in mind. We sprinkle comments and
 445 remarks throughout the text, in the hope that it provides useful guidance
 446 with respect to the big picture. The book assumes the reader to have math-
 447 ematical knowledge commonly covered in high-school mathematics and
 448 physics. For example, the reader should have seen derivatives and inte-
 449 grals before, and geometric vectors in two or three dimensions. Starting
 450 from there we generalize these concepts. Therefore, the target audience
 451 of the book includes undergraduate university students, evening learners
 452 who and people who participate in online machine learning courses.

453 In analogy to music, there are three types of interaction, which people
 454 have with machine learning:

455 *Astute listener*

456 The democratization of machine learning by the provision of open-source
 457 software, online tutorials, and cloud-based tools allows users to not worry
 458 about the nitty gritty details of pipelines. Users can focus on extracting

459 insights from data using off-the shelf tools. This enables non-tech savvy
460 domain experts to benefit from machine learning. This is similar to lis-
461 tening to music; the user is able to choose and discern between different
462 types of machine learning, and benefits from it. More experienced users
463 are like music critics, asking important questions about the application of
464 machine learning in society such as ethics, fairness, and privacy of the in-
465 dividual. We hope that this book provides a framework for thinking about
466 the certification and risk management of machine learning systems, and
467 allow them to use their domain expertise to build better machine learning
468 systems.

469 *Experienced artist*

470 Skilled practitioners of machine learning are able to plug and play differ-
471 ent tools and libraries into an analysis pipeline. The stereotypical prac-
472 titioner would be a data scientist or engineer who understands machine
473 learning interfaces and their use cases, and is able to perform wonderful
474 feats of prediction from data. This is similar to virtuosos playing music,
475 where highly skilled practitioners can bring existing instruments to life,
476 and bring enjoyment to their audience. Using the mathematics presented
477 here as a primer, practitioners would be able to understand the benefits
478 and limits of their favorite method, and to extend and generalize existing
479 machine learning algorithms. We hope that this book provides the impe-
480 tus for more rigorous and principled development of machine learning
481 methods.

482 *Fledgling composer*

483 As machine learning is applied to new domains, developers of machine
484 learning need to develop new methods and extend existing algorithms.
485 They are often researchers who need to understand the mathematical ba-
486 sis of machine learning and uncover relationships between different tasks.
487 This is similar to composers of music who, within the rules and structure
488 of musical theory, create new and amazing pieces. We hope this book pro-
489 vides a high-level overview of other technical books for people who want
490 to become composers of machine learning. There is a great need in society
491 for new researchers who are able to propose and explore novel approaches
492 for attacking the many challenges of learning from data.