Machine learning: The Complete Beginnerâ€™s Guide to Learn and Effectively Understand Machine Learning Techniques (Intermediate, Advanced, To Expert Concepts)

Pages: 122
Format: pdf, epub
Language: English

MACHINE LEARNING

The Absolute Beginner’s Guide to Learn and Effectively Understand Machine Learning From Beginners, Applications of Deep Learning (Intermediate, Advanced, To Expert Concepts)

A

ANTONIO ROBERT

ASIN: B07TXLXDCW

ISBN: 9781077867420

TABLE OF CONTENTS

CHAPTER ONE:

INTRODUCTION TO MACHINE LEARNING

The machine learning (English learning machine, literally the " machine learning ") or machine learning is a field of study of artificial intelligence that is based on statistical approaches to give computers the ability to "learn" from data, that is to say to improve their performance to solve tasks without being explicitly programmed for each. More broadly, this concerns the design, analysis, development, and implementation of such methods.

Machine learning usually has two phases. The first is to estimate a model from data, called observations that are available and infinite numbers, during the design phase of the system. Model estimation involves solving a practical task, such as translating a speech, estimate a probability density, recognize the presence of a cat in a photograph, or participate in the driving of an autonomous vehicle. This so-called "learning" or "training" phase is generally performed before the practical use of the model.

The second phase corresponds to the setting in production: the model is determined, new data can then be submitted to obtain the result corresponding to the desired task. In practice, some systems can continue their learning once in production, provided they have a way to get a return on the quality of the results produced.

According to the information available during the learning phase, learning is qualified in different ways. If the data is tagged (that is, the response to the task is known for that data), it is supervised learning. There are classifications, If the labels are discrete, or of regression or if they are consecutive. If the model is learned incrementally based on a reward received by the program for each of the actions taken, it is called reinforcement learning. In the most general, unlabeled case, one seeks to determine the underlying structure of the data (which may be a probability density), and it is unsupervised learning. Machine learning can be applied to different types of data, such as graphs, trees, curves, or simply feature vectors, which can be continuous or discrete.

A LITTLE HISTORY OF MACHINE LEARNING

Machine Learning (ML) emerged in the second half of the 20th century from the field of artificial intelligence and corresponded to the development of algorithms capable of accumulating knowledge and intelligence from experiences, without being humanly guided during their learning, nor explicitly programmed to manage such or such specific experience or data.

The Beginnings Of Artificial Intelligence

Artificial intelligence (AI) is one of the newest and most successful scientific and engineering disciplines. The pioneers of AI began their work towards the end of the Second World War. In 1956, ten pioneers of American research in the theory of automata, neural networks, and intelligence came together for a two-month workshop at Dartmouth College: J. McCarthy (Dartmouth), Minsky (Princeton), C. Shannon (Bell Labs / MIT), N. Rochester (IBM), T. More (Princeton), A. Newell (Carnegie Tech), H. Simon (Carnegie Tech), A. Samuel (IBM), R. Solomonoff (MIT)), O. Selfridge (MIT). It was throughout this workshop that the term AI appeared: the AI was officially born.

Quickly, the AI

sparked a lot of debate among specialists as to the approach to adopt. Several schools have emerged. Behaviorists sought to build AIs that could think or act like humans. Rationalists turned to a rational view of AIs that they thought should produce rational workings and rationalizations, rather than trying at all costs to imitate the human brain.

Rationalists therefore first tried to build programs that could solve any problem described by a logical notation. Nevertheless, the translation of a complex or informal knowledge within the logic is very difficult. On the other hand, building AI capable of producing reasoning and rational behavior assumes the exact processing of many parameters: a computing capacity that was inaccessible at the time.

In the behaviorist approach, the Turing test (very recently back on the scene with Eugene Goostman, or some high-profile films like Ex Machina was designed to define intelligence. IA passes the Turing test if a human is not able to determine in a conversation whether the AI

is a human or a computer. This feat requires several capabilities: automatic processing of human language (Natural Language Processing), representation (ordered storage) of knowledge and past experiences, faculties of automatic reasoning and ML, that is to say, a capacity of adaptation according to the circumstances, this set of faculties necessary for intelligence are particularly difficult to develop simultaneously in an AI.

The Singularity

The concept of technological singularity contains the idea that, from the point of human evolution (the singularity), AIs will go beyond the human performance of reasoning on the one hand and the evolving capacities of humanity, on the other hand. AI would then carry progress that human society would no longer be able to control. The idea dated back to the 1950s with John von Neumann and was debated by IJ Good, V. Vinge or R. Kurzweil. It is frequently found in popular culture with 2001 Space Odyssey, Terminator, Minority Report, Matrix, and more recently Her or Transcendence.

The relevance of this concept is still highly contested and debated in the scientific community today. We are content here to mention the existence of the concept because it is interesting to know it, but without engaging in debate.

First Advances

McCulloch and Pitts conducted the earliest recognized AI work in 1943 with the first model of artificial neurons: each neuron is characterized as open or closed and can open in response to stimulation by a neuron, a sufficient number of neighboring neurons to transmit a signal. Another pioneer of AI, Hebb, proposed in 1949 a system of neurons where the connections between neurons are weighted with rules of modification of the weights constituting learning.

The period that followed until the 1960s was punctuated by great advances and great hopes: the "General Problem Solver" of Newell and Simon, who tried to imitate the human methods of solving problems; Gelernter's "Geometry Theorem Prover", capable of proving difficult mathematical theorems; the first chess programs (A. Samuel); the creation of Lisp (J. McCarthy), which later became the dominant programming language for AIs; McCarthy's "Advice Taker", a hypothetical program built to use the knowledge acquired to solve problems: the first appearance of the principles of formal and explicit representation of knowledge and the manipulation of this representation by deduction processes.

Many people, including government program managers, have become inflamed by imagining AIs capable of reaching the capabilities of the human brain.

Stopping Ai

After the rush of the 1950s, the 1960s brought more realism and pragmatism to the field. First, it was realized that most AIs function mainly by simple syntactic manipulations and not by truly knowledge-based reasoning. In the case of translation, for example, the machine needs to have a fine understanding of the context to translate without ambiguity. AIs failed to translate even simple texts. In addition to translation, many problems (mathematical theorems, deductive reasoning, and complex problems involving several sub-tasks) proved impossible to solve because of the number of parameters involved.

The proposed neural network structures were far too limited in size and the ability to represent knowledge to hope to produce intelligent behavior. In 1969, Minsky and Papert demonstrated that perceptrons (simple models based on the work of McCulloch and Pitts) could only construct a very minimal representation: for example, they were not able to implement the XOR function.) which allows recognizing if two inputs are different.

The prospect of building more complex systems was out of step with the amount of data available and the storage (that is, the representation of knowledge) needed. On the other hand, the computing power of computers of the time did not allow implementing the learning of these models. The hopes placed in AI were therefore considerably disappointed, and many research programs were stopped. This has pointed to a sharp slowdown in research investment, especially that focused on neural networks and seeking to replicate human reasoning.

Pragmatism And Modestice: The Emergence Of The Learning Machine

Between the 1970s and 1980s, efforts were focused on systems of expertise restricted to specific areas such as chemical analysis, medical diagnostics, industry-leading systems, and highly specialized robots or software.

Work on Markov models, theoretical and technical advances in Monte Carlo simulations and the development of Bayesian network formalism have punctuated this period of twenty years, with use in specific fields with relatively few interdisciplinary exchanges. The late eighties also saw the return of neural networks, with the (re) invention of backpropagation learning and its application to specific problems such as character recognition. But it will be necessary to wait until the mid-2000s to see successes on more complex problems and in wider domains.

The emergence of ML as a discipline in its own right is part of this dynamic of pragmatism and modesty; it took place in the 1990s with the desire to focus on solving very concrete and humble problems, in direct relation to the available data.

At the equivalent time, the advent of the Internet has led to stronger collaborations between different research groups, faster communication of results and, above all, the ability to replicate experiences using test data and code repositories. Interdisciplinary with engineering and computer science has also made it possible to advance much faster on these systems. With a more systematic scientific approach, the experts were able to concentrate their efforts on improving existing and promising theories and models together, rather than operating in isolation.

It should be noted that it is at this point in history that mathematics and especially statistics have been reintroduced into the development of the theoretical and applied aspects of ML.

Input Of Big Data

It is also critical to note the significance of the arrival of Big Data: the availability of more varied sets of massive data and the increase of computing and storage capacities.

The development of Big Data technologies has made it possible to overcome important limits of the ML: the absence of a sufficient amount of data does not allow driving a prediction algorithm effectively. Therefore, in the absence of such data, there is no convincing experimentation; it is then more difficult to develop the theoretical foundations of an ML algorithm.

Big Data also poses theoretical challenges for ML, such as large-scale problems: the cross-fertilization of many data sources greatly dilutes information. This increases the number of variables to explore, but not the number of examples available.

What Is Ml?

The ML is a sub-part of artificial intelligence that focuses on creating machines that behave and operate intelligently or simulate that intelligence.

AI does not necessarily mean learning the machine. For example, in the case of machines that play chess: An AI will not be able to cope with a new strategy of a human player by learning from the experience to adapt his behavior unless it was conceived within the framework of ML.

For a machine to automatically learn and adapt, the knowledge it can extract from its experiences must be stored in some form, so that it can be used for a specific purpose. In general, the goal is to predict the future (a behavior, a number, etc.), but a machine can also be constructed to artificially generate new data (generative machines) statistically similar to the original data or to detect patterns functioning (causalities, the structure of a network, etc.).

As mentioned above, ML has drawn heavily from mathematics and statistics. Classically, the statistics aim to allow a global understanding of the data: what are the characteristics of a variable (minimum, maximum, average)? More advanced, what is the statistical distribution? Can we model the data generation process by trying to extract the noise contained in the data? The simplest model is regression, which is widely used in many fields. To train (to teach) a model of regression is to make ML. We can, therefore, trace the machine learning in the seventeenth century with Legendre and Gauss and their method of least squares.

The ML can be seen as the set of methods to determine the best way to model the data. Of these methods, the most central and most important is learning. Learning consists, for a given model, of choosing the best possible parameters to describe the data. For a simple regression, it is a question of choosing the best pair a, b. In this case, the knowledge acquired by the experiment (the different data samples are as many micro-experiments), is stored in these two coefficients a and b. A more complex model stores its knowledge in a much larger number of coefficients that can exceed a million or even a billion (for complex neural networks for example), but the principle is the same: the ML allows a machine (the model,

Models and their associated learning procedures include ideas derived directly from statistics and mathematics, but not only. Pragmatism has led research to push back interdisciplinary boundaries. Physics (including statistical physics for Boltzmann machines, a special case of neural networks), biology and neurobiology, imaging and engineering have been very important sources of ideas for developing ML models and learning techniques.

They group the ML algorithms into two or three major classes, which correspond to different types of learning:

Supervised learning: we give the algorithm a certain number of examples (input) to learn from, and these examples are "labeled," that is, we associate them with the desired result (output). The algorithm then has for the task to find the law, which makes it possible to find the output according to the inputs.

A typical example is the classification of email into two classes: "spam" or "ham." The input, in this case, is the email, either raw or formatted (prepared as a set of "features," or variables) to optimize the performance of the algorithm. The output is a label, 0/1, or Spam / Ham.

It is also possible to establish medical diagnoses: whether or not a patient has cancer (output), according to his symptoms and phenotypes (input)

Another example: the recognition of handwritten digits (the image is a "1" or a "9"?). In this case, the input is a set of pixels, and the output takes 10 values, from 0 to 9.

Unsupervised learning: here, no label is provided to the algorithm, and it must discover without human assistance the characteristic structure of the input.

The typical example is clustering. The algorithm will group the examples into different clusters or classes. A use case is the recognition of images: a clustering algorithm will be able to group images of cars in the same cluster, and images of buildings in another cluster. He will be able to do this because he knows how to learn that a car has the characteristic of having wheels, and a particular shape and a building of more straight forms. This learning, the algorithm does it alone, without a human explicitly introduces the concept of car or building, nor explains to him that he must seek the presence of wheels to identify a car. He alone learns the concept of the curve, right, corner, etc., independently of human representations. This independence can sometimes make the functioning of the algorithm quite opaque and difficult to interpret.

The human being is led by his fellow beings to represent the world and to interpret his experiences like the others; it is for this reason that we can communicate between us descriptions. Our human representations are very oriented towards the view. The other senses have a representation that is perhaps less consensual. Two different people will have some difficulty in explaining each other's taste experiences by default of common representation. It's a bit the same with an image recognition algorithm. He will build his representation, and a human may have difficulty understanding it. The human being is led by his fellow beings to represent the world and to interpret his experiences like the others; it is for this reason that we can communicate between us descriptions. Our human representations are very oriented towards the view. The other senses have a representation that is perhaps less consensual. Two different people will have some difficulty in explaining each other's taste experiences by default of common representation. It's a bit the same with an image recognition algorithm. He will build his representation, and a human may have difficulty understanding it. The human being is led by his fellow beings to represent the world and to interpret his experiences like the others; it is for this reason that we can communicate between us descriptions. Our human representations are very oriented towards the view. The other senses have a representation that is perhaps less consensual. Two different people will have some difficulty in explaining each other's taste experiences by default of common representation. It's a bit the same with an image recognition algorithm. He will build his representation, and a human may have difficulty understanding it. The other senses have a representation that is perhaps less consensual. Two different people will have some difficulty in explaining each other's taste experiences by default of common representation. It's a bit the same with an image recognition algorithm. He will build his representation, and a human may have difficulty understanding it. The other senses have a representation that is perhaps less consensual. Two different people will have some difficulty in explaining each other's taste experiences by default of common representation. It's a bit the same with an image recognition algorithm. He will build his representation, and a human may have difficulty understanding it.

Reinforcement learning: it's a bit of an intermediary between the first two types of learning. The algorithm is designed to seek to optimize at any cost a quantitative reward, positive or negative, from experiences corresponding to different situations, an example of the most current and impressive learning reinforcement is the Boston Dynamics Big Dog robot. It is a four-legged robot, able to walk, run, climb, and carry heavy loads. The robot is trained to walk by reinforcement: its reward is strong and positive if it stays up, its reward is negative if it falls. He is programmed to explore the different movements he can perform and make his own experiences. This is how he learns alone to climb a slope and optimize the way to put his legs (speed, frequency, angle, etc.) in a way very similar to an animal or a human.

Supervised learning is a little more intuitive, in part because many supervised learning algorithms have been used successfully in many cases. Among the most popular and widely used supervised learning models are linear regressions, logistic regression.

In 1984, Leo Breiman, Jerome Friedman, Richard Olshen and Charles Stones pose one of the cornerstones of the ML, introducing the theoretical foundations of modern decision tree learning, with innovative techniques and algorithms for processing large quantities of data. Their work was very influential, and in 2001, Leo Breiman and Adele Cutler introduced random forests, an extremely popular and powerful model of ML today. Similarly, the boosting gradient, another very popular type of model is the product of the work of Breiman, Friedman, Mason, Baxter, Bartlett, and Frean.

In recent years, neural networks have returned to the spotlight. They are a separate domain of machine learning. As we have seen, their story began in the 1950s, and they suffered several setbacks until 1989, with the successful application of the backpropagation algorithm to deep neural networks (LeCun), and in 2006 with several important technical advances in the learning of neural networks (in the works of Hinton, Bengio, Ranzato, LeCun, Larochelle).

Trees and decision forests, neural networks are just a few popular examples of ML models. Many other models exist and are used successfully: Bayesian learning, Bayesian networks, machine vector support, k-means (1992), DBSCAN (1996), OPTICS (1999) for clustering, or even genetic algorithms whose foundations go back to 1975.

Good knowledge of theoretical foundations and learning techniques makes it possible to choose the model best suited to modeling. This is one of the roles of the data scientist in data analysis projects.

Reunification

Over the last 20 years, AI has focused on designing intelligent agents that can perceive and act in a given environment. Machine learning algorithms have developed and become more and more complex. The border between the two domains is now disappearing. To solve complex difficulties involving various phenomena at different scales, research tends to combine IA, ML, robotics, and algorithmics in powerful and complex systems.

Boston Dynamics robots like Big Dog or WildCat are impressive examples: they combine robotics engineering with intelligent machines driven with ML to learn from experience.

The Challenges Of Ml

In recent years, the arrival of Big Data and the computing power available have led to great advances in ML, including optimized learning techniques for models as simple as regressions. The challenges are still there today for research. We must still lay the theoretical foundations of entire ML domains: indeed, some models are known and effective, but the knowledge is sometimes very phenomenal and based on the observation of many experiments.

Machine learning is one of the principal areas of artificial intelligence.

It concerns the study and the development of quantitative models allowing a computer to accomplish tasks without it being explicitly programmed to do them. Learning in this context means recognizing complex shapes and making smart decisions. Given all the existing entries, the complexity of doing so lies in the fact that the set of possible decisions is usually very difficult to enumerate. The machine learning algorithms have, therefore, been designed to gain knowledge about the problem to be addressed based on a set of limited data from this problem.

This guidebook is going to take some time to explore machine learning, and what it is all about. There are so many different aspects of machine learning and how to make it work for your needs, and all of it is found in this guidebook. Some of the different topics that you will be able to learn about inside include:
Get access to free software and data sets so you can try out your very own machine learning software. See how advanced machine learning will impact our world in the future!

Also, this book presents the scientific foundations of the theory of supervised learning, the most widespread algorithms developed in this field as well as the two frameworks of semi-supervised learning and scheduling, at a level accessible to master's students and engineering students. We had here the concern to provide a coherent presentation linking the theory to the algorithms developed in this sphere. But this study is not limited to present these foundations; you will find some programs of classical algorithms proposed in this manuscript, written in C language (language both simple and popular), and for readers who want to know how it works. These models are sometimes referred to as black boxes.

Who is this book directed to:Â Â

The engineering students, master’s students, including doctoral students in applied mathematics, algorithmic, operations research, production management, decision support.

Also, to engineers, teacher-researchers, computer scientists, industrialists, economists, and decision-makers who have to resolve problems of classification, partitioning, and scheduling on a large scale.

In this book, you will attain helpful information for getting started, such as:

Why Use Neural Networks?
The Different Types Of Learning
Machine Learning In PracticeÂ Â Â Â
Reinforcement LearningÂ
Learning by reinforcementÂ Â Â Â Â Â Â Â Â
Neural NetworksÂ Â Â
Tasks Of Neural Networks
Neural Networks versus Conventional ComputersÂ Â
Deep LearningÂ Â Â Â Â Â Â Â
Machine Learning and Data MiningÂ Â Â Â Â
Running Python Getting Started and more,

When you are ready to learn more about what machine learning is all about, and how you are able to benefit from it in your own coding and programming, make sure to check out this guidebook to help you get started.

Do not waste time to gather partial or false information, when you can get everything you require to REACH YOUR GOALS by reading this fantastic guide.

Scroll Up and Click the Buy Now Button!

<

Machine Learning Week 1 Quiz 1 - This course will cover Chapters 1-5 of the textbook â€œPython for Everybodyâ€•. This course will be a quick way to understand all the major concepts of Python Learn Python for AI, Machine Learning, Data Science and App Development Complete Python Programming Masterclass Beginner to Advanced Udemy Free Start Here with Machine Learning - Machine Learning Mastery - Learning Python written by Python expert Mark Lutz will give you comprehensive and The book is a complete treat for the beginners who want learn a programming Learn Python Programming: The no-nonsense, beginner's guide to. many concepts, models, advanced methods, machine learning algorithms and many The Best Python Books for All Skill Levels - Stack Abuse - Whereas with Udacity they actually teach you all the concepts you need to solve For intermediate students, advanced machine learning algorithms, big data, deep [Udacity] AI Programming With Python Nanodegree Learn Python, NumPy, coming from calculus and statistics classes â€¢ Answered homework, textbook, Self Taught Python Programmer - Get this book free when you sign up for a 30-day Trial. Machine Learning: 2 Manuscripts in 1 Book: Machine Learning for Beginners and Machine Learning.. you need to know about machine learning and has the following guides: This audiobook was very good in terms of projects and concept understanding. I got to The Best Python Books â€“ Real Python - If you wish to learn Python programming, there are plenty of books This book is also for programmers who have a vague understanding of. Introduction to Machine Learning with Python: A Guide for Data Scientists You also learn about the advanced methods for model evaluation and the concept of Machine Learning Week 1 Quiz 1 - Machine Learning â€” Andrew Ng. If you are a Python or a Ruby programmer, this. that I talked about, and said that, you know Is my comprehension correct? Guide to Machine Learning, if you want to read more about the concepts of It starts by covering linear regression and progresses towards higher level algorithms. Machine learning: The Complete Beginner's Guide to Learn - Complete the quiz and submit it through Blackboard by the end of the day on This 29-part course consists of tutorials on ML concepts and algorithms, as well Learn Python Machine learning by

Quiz (Beginners to Experts) Learn Python effective web search, and a vastly improved understanding of the human genome. Learning Python â€” The Hitchhiker's Guide to Python - Go from zero to hero in Machine Learning with this complete course. The best way to learn Python is to understand the big picture of all what you need to you are a complete beginner and by the end of the course you will be at advanced level. Expert in Machine Learning and author of the Zero to Deep Learning book. Learning golang from zero to hero - Innovators of Change - Learn Neural Networks and Deep Learning from deeplearning. a complete, stand alone interpretation of Stanford's machine learning course presented. Learning Algorithms in Python and R from two Data Science experts.. This is an intermediate-level specialization, and it is assumed you already know how to write Meta Learning Tutorial - Whether you're looking advance your Python learning or just beginning to dip 80 data structures and algorithm exercises and guidance from a dedicated mentor to In this book, you will learn the essential Python data structures and the most. deep understanding of Data structures and algorithms concepts from ground Machine Learning Week 1 Quiz 1 - If you are currently enrolled in or completed Come learn Machine Learning with Meanwhile the courses he teaches at intermediate to advanced level have just Learning Algorithms in Python and R from two Data Science experts. effective web search, and a vastly improved understanding of the human genome.
