By Paper.li-4.02.2021-The Garrison Platoon Of Books:

How To Read 43 Machine Learning Books in a Year

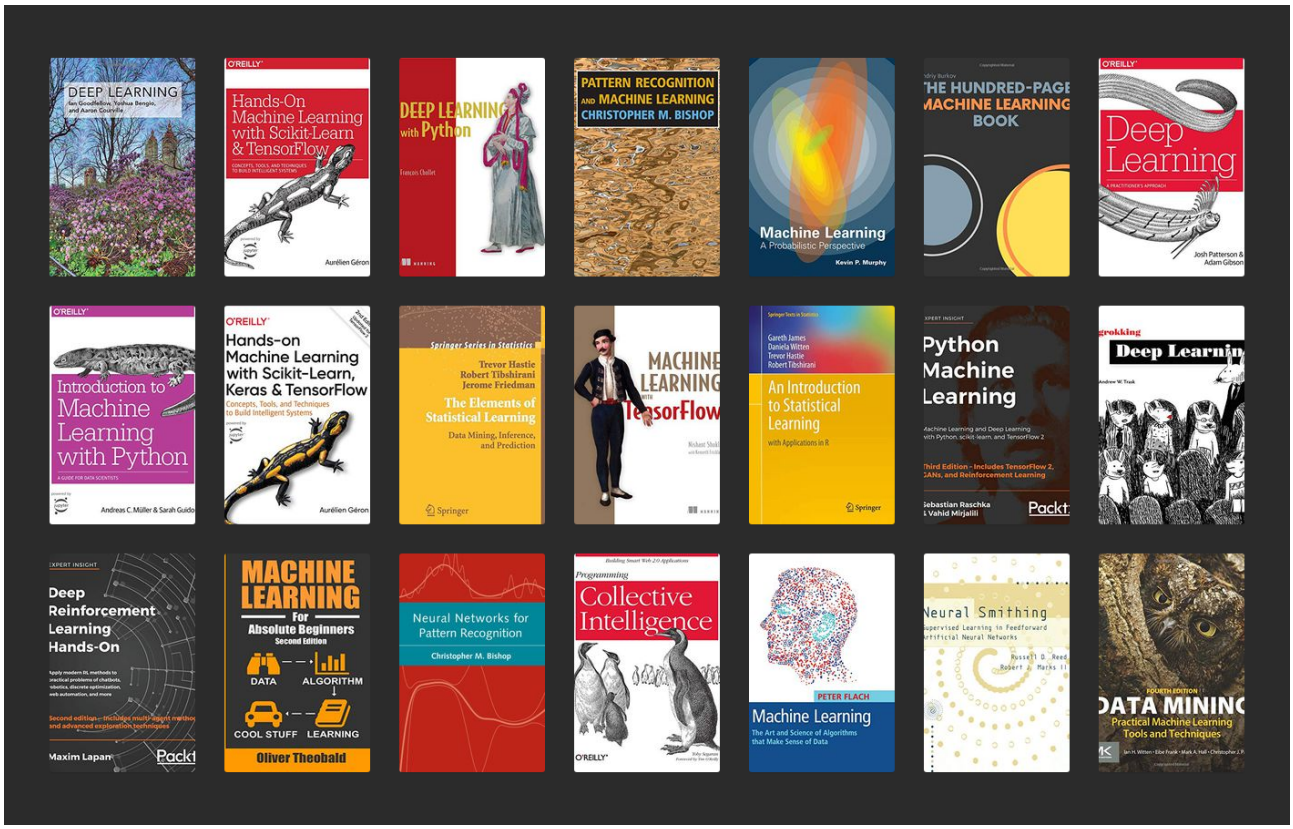

GANAPATHI PULIPAKA

Dr Ganapathi Pulipaka is Chief AI HPC Scientist and bestselling…

The Garrison (armies of militia) of libraries worldwide offer millions of books, such as the Library of Congress in D.C. has over 162 million books, and the New York Public library carries around 53 million books. So many books, so little time in a human's life.

A number of people have asked me through several of my channels and conferences — how to find time to read books, and what can be done to read more books each month. Some audiences even feel that 43 machine learning books in a year are insufficient, and want more.

I keep discovering new material every day on top of the antiquated books, which still offer good concepts. To get started, I would suggest disconnecting from Netflix, Amazon Video, and regular TV channels. The more you watch any of this stuff, the more you wouldn't be finding time to read the books.

In 2020, I had managed to read more than 96,120 pieces of books, eBooks, articles, averaging 267 pieces of books, eBooks, research papers, or articles per day. However, on average, people might have read 10 to 30 machine learning books in a year.

Reading through machine learning literature for machine learning and deep learning algorithms, high-performance computing, cloud computing, IoT, and IIoT usually does not offer me something new. Often, we have newspapers, new techniques, or algorithms that are released well over the month. But, the audience usually associates their memories of recognizing each book through photographic memory.

Machine learning engineering and data engineering

There are many books surrounding machine learning topics, including ways of implementing algorithms with PyTorch, C, C++, R, TensorFlow, and Python. People with prior experience in Python and Python's scientific libraries like Pandas, Numpy, and Matplotlib, do not feel the challenge to work on TensorFlow or PyTorch. Primarily, getting started would require understanding and fitting the machine learning model to data by optimizing the cost function, the bias, variance, and trade-off. However, data engineer readers are more focused on data pre-processing, cleaning, data wrangling, and creating balanced datasets. In this case, they should

concentrate more on those chapters and sections that provide techniques and tips to expedite their needs and provide an understanding of advanced techniques.

Another aspect to look at machine learning books would be to develop skills in neural networks and deep learning with a number of neural network architectures (such as convolutional neural networks, recurrent neural networks, LSTM, multilayer perceptron, backpropagation, and activation functions) along with feedforward architecture in neural networks for weight summation calculation process and distributed deep learning for large-scale datasets.

Reading books can get you prepared to understand the concepts of mathematics and statistics before you can take a course with Quickstart (where I'm currently teaching), Coursera, or Columbia Engineering University for gaining your certifications. As a matter of fact, machine learning and data science certifications require few prerequisites that can help candidates to reduce the time you spend on the course.

Reading a book also depends on what you're trying to gain. For example, if you're looking for semi-supervised machine learning, then you need to jump to the section of that chapter. If you're looking for supervised, unsupervised, or reinforcement learning examples, then you need to concentrate on those sections. On the other hand, if you're already familiar with the machine learning

algorithm, just jump to the exercises section in a book and start implementing a linear regression project, logistic regression, support vector machines, decision trees, random forests, and ensemble learning.

Suppose you do not have prior experience in any programming language and want to learn a machine learning programming language such as Python; in that case, you can start with the Python Crash coursebook. With this, one can begin building momentum towards understanding what type of data variables are available in Python, string operations, number operations, working with lists and dictionaries, and can work on few projects for data visualization, generating data, plotting data points, API calls, integrating and developing web applications with Django, and deploying apps.

Similarly, one can work their way up on PyTorch, TensorFlow, and R. Once you're comfortable dealing with primary operations in PyTorch, R, TensorFlow, and Python, you can work on the machine learning algorithms. Here, one needs to find the balance between mathematics, statistics, and programming languages. Once completed the process of gaining expertise in mathematics, statistics, PyTorch, TensorFlow, and Python, one can quickly read the books on machine learning. One should be able to assess the model accuracy, quality fit (avoiding overfitting, underfitting), bias,

variance, trade-off, and implement all the machine learning algorithms.

My upcoming book covers machine learning algorithms for laying a greater foundation on mathematics with linear algebra, calculus, statistics, probability theory, natural language processing, bioinspired machine learning algorithms, deep learning algorithms, supervised and unsupervised machine learning, and implementation of reinforcement learning algorithms in PyTorch, TensorFlow, R, and Python. With the book, one can go for a vertical take-off from the runway to launching their data science career.

Rapid reading techniques

SEE ALSO

OPINIONS

Case Study: How This Women-Owned Milk Producer Company Enabled A Smoother Payment Process With SAP-ERP

Reading machine learning books should be more like riding a bike or driving a car. The fundamental building blocks for reading many books come from concentration and focus. If you're a rapid reader and read 2000 words per minute, you will lose tons of sentences, paragraphs, concepts, and the comprehension significantly goes down every second you get distracted. You can easily break the barriers of 2000 wpm in a month. It's like driving a car and looking at the road, but with millions of thoughts on your mind, you're going to miss the exit and can end up going to a different city. Make sure that you're not running into constant interruptions. In fact, climbing a mountain with a phone in your hand will only slow you down or bring you down unless you're hanging on a rope to survive.

The brain functions at delta, theta, alpha, beta, and gamma frequencies, and thus, the state of mind needs to associate with alpha and beta frequency. There are several other memory techniques each person can adapt based on their brain frequency.

Rapid readers can either leverage traditional speed reading or subconscious photographic-type reading. As I mentioned above, subconscious reading helps to get that layered reading required to

increase their awareness and familiarity with a topic. Anytime you're having regression (a word used in rapid reading to describe the lapse in concentration that makes you read the same paragraph many times repeatedly), it can break your flow of reading.

The other great technique I use is skimming. When I read a paragraph, I read with higher comprehension, move rhythmically, and skim the paragraph, laser-focusing on the topic's concept in the section. One should read the page through a vertical line in the middle of the page's text, rather than reading from the beginning of the sentence and guiding yourself to read line by line.

One should also not focus on the hand-movement while scrolling through the pages. I strongly suggest minimizing sub-vocalization of reading word by word as one continues reading many pages and chapters. Let the familiarity of the topic be the light of your guide that will allow you to skim the pages. The subvocalization would not let the readers get past 600 wpm.

However, I find speed reading, that I picked up during my childhood through a vertical line rapid reading technique, the most effective reading technique. With this technique, I have managed to complete reading each book in just a few minutes.

What Do You Think?

Download our Mobile App. WHAT'S YOUR REACTION?

EXCITED
0

HAPPY
0

IN LOVE
0

NOT SURE
0

SILLY
0

GANAPATHI PULIPAKA

Dr Ganapathi Pulipaka is Chief AI HPC Scientist and bestselling author of books covering AI infrastructure, supercomputing, high-performance computing for HPC, parallel computing, neural network architecture, data science, machine learning, and deep learning in C, C++, Java, Python, R, TensorFlow, and PyTorch on Linux, macOS, and Windows.

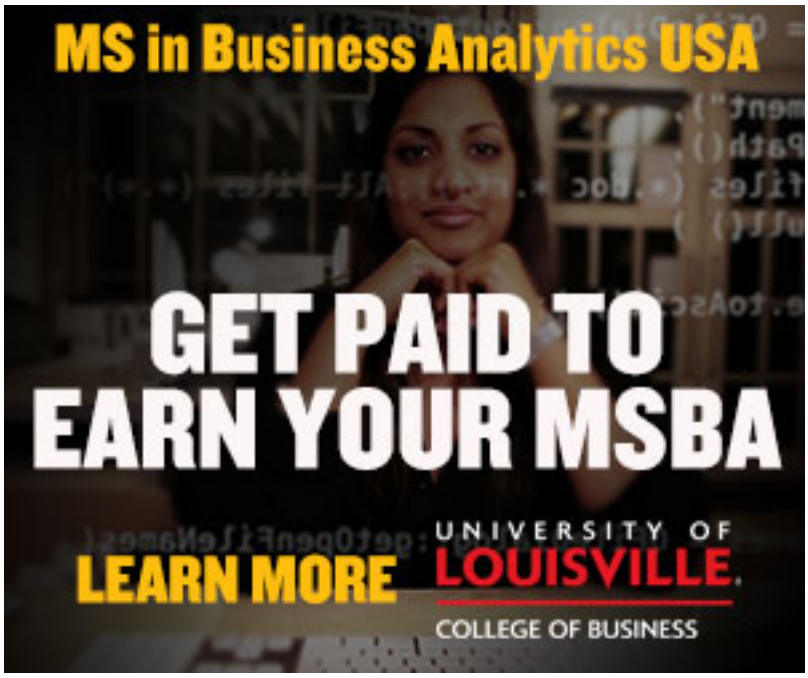

OUR UPCOMING EVENTS

- Workshop I Hands-on Deep Learning Model Deployment & Management I Register here>>
- Machine Learning Developers Summit 2021 I 11-13th Feb I

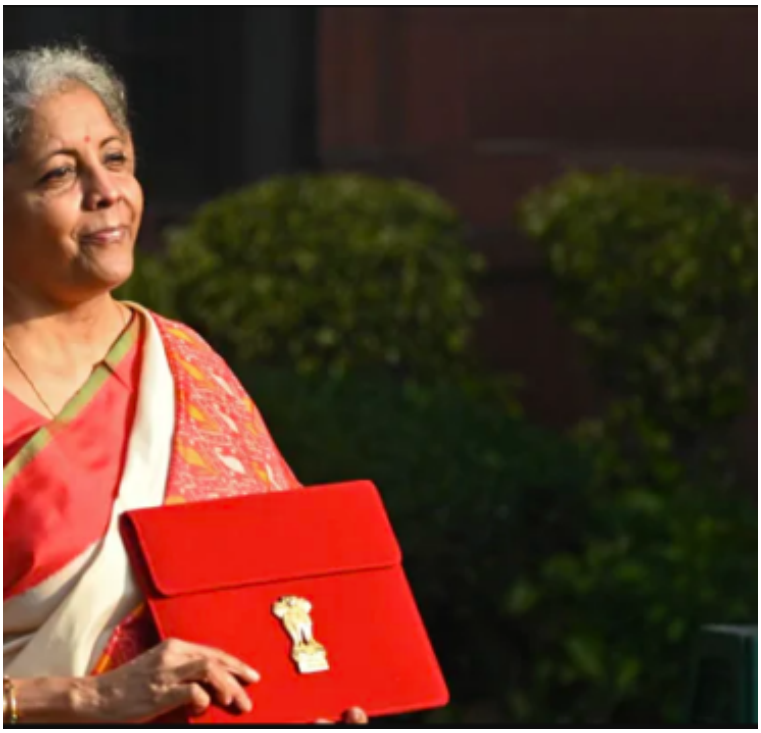

Machine Learning 101: Ten Projects For Beginners To Get Started

- 28/01/2021
- 4 MINS READ

NEWS

DRDO Announces Online Courses On Artificial Intelligence, Machine Learning & Cybersecurity

- 27/01/2021
- 2 MINS READ

STARTUPS

Enabling Connected Vehicle Ecosystem With Machine Learning: The Startup Story Of Sibros

- 27/01/2021
- 4 MINS READ

DEVELOPERS CORNER

Complete Guide To AutoGL -The Latest AutoML Framework For Graph Datasets

- 20/01/2021
- 5 MINS READ

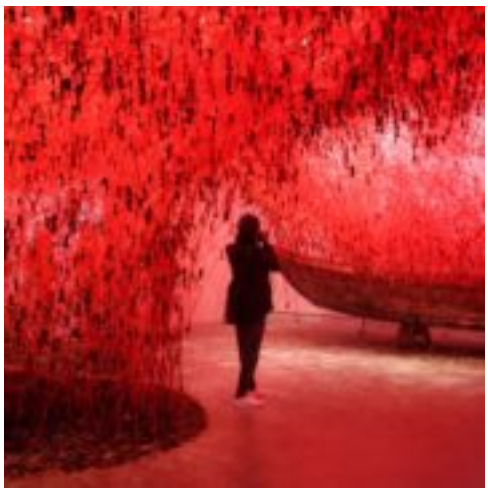

DEVELOPERS CORNER

Guide To MLflow – A Platform To Manage Machine Learning Lifecycle

- 20/01/2021
- 7 MINS READ

EVENTS

MLDS 2021: Top Talks You Should Definitely Attend This Year

- 19/01/2021
- 3 MINS READ