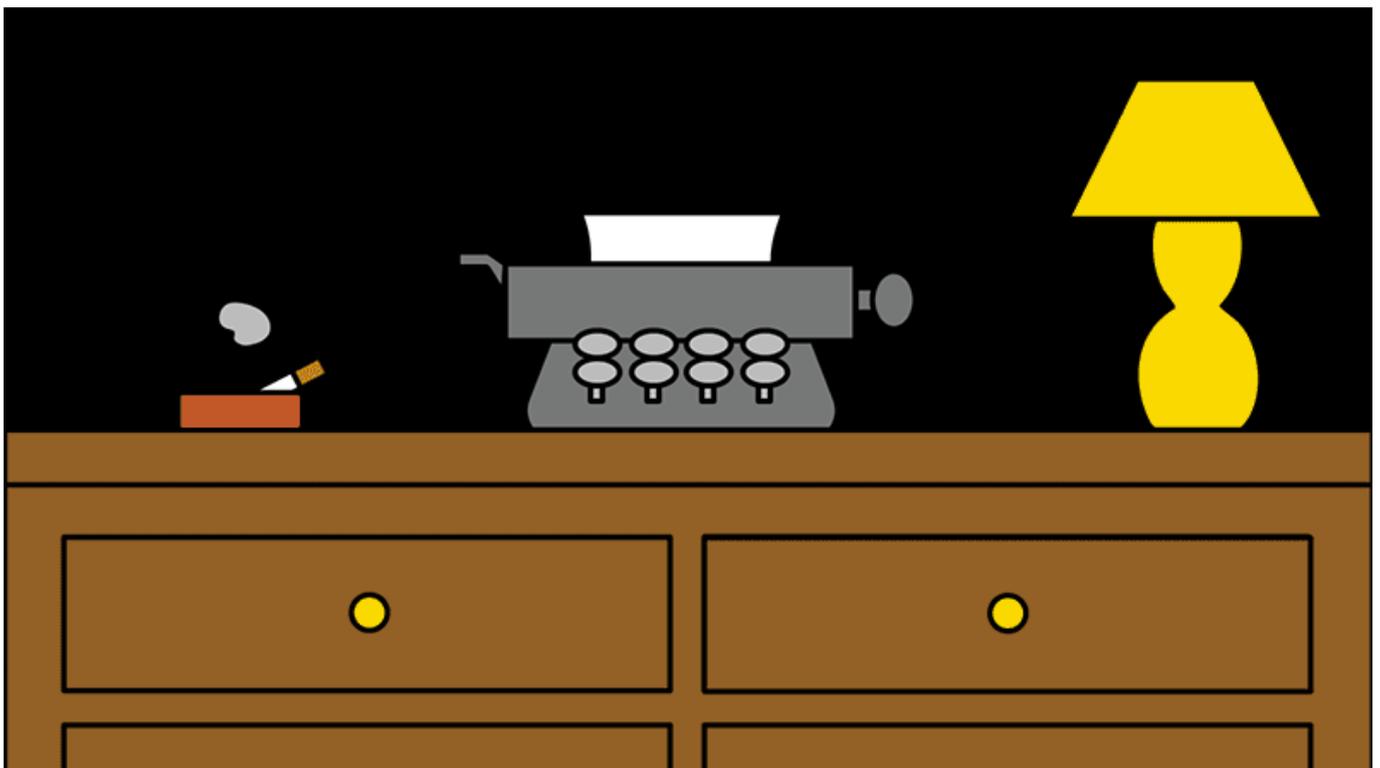


5 Skills You Need to Become a Machine Learning Engineer

Machine learning Skills are exceedingly high in demand for it has resulted in a plethora of innovations and improvements among businesses and organizations worldwide.

It has also provoked reactions ranging from curiosity to anxiety among people.

There are many facets to Machine Learning and given how rapidly the Machine Learning space is evolving, I started creating a review-driven guide that highlights the Machine Learning skills that companies are actively seeking.



I've put a tremendous amount of efforts in consulting Data Scientists, Machine Learning Engineers, Educators and a few companies to identify the core skills for learners to be successful.

These suggestions are derived from conversations, interviews and collaboration with Educators, as well as my own experience in both Machine Learning and industry roles.

These skills are taught excellently in [Machine Learning Career Track](#) with Unlimited 1:1 Coaching and Job Guarantee.

Machine Learning offers great potential and in this guide, we interrogate what's possible.

Mindset and Skills You Need to Become a Machine Learning Scientist/ Engineer

This piece will help you to understand the mindset and the specific skills you'll need to prepare for Machine Learning roles .

— Sound Understanding of the Ecosystem

Before getting into specific skills required for machine learning roles, there is one concept i want to address.

Being a Machine Learning Scientist/ Engineer heavily depends on having a crisp awareness about the entire ecosystem that you're designing for.

Let's say that you are working for an App based food delivery company, and your company wants to issue targeted coupons based on things like preferences of customer, order history, with a goal of generating coupons that a customer will use.



In your Data Analysis model, you can collect the order history, purchase data, do the analysis to figure out behaviour, preferences, trends, and then propose strategies that might work.

However, The Machine Learning approach would be simple, i.e; use your data analysis model and write an automated coupon generation system.

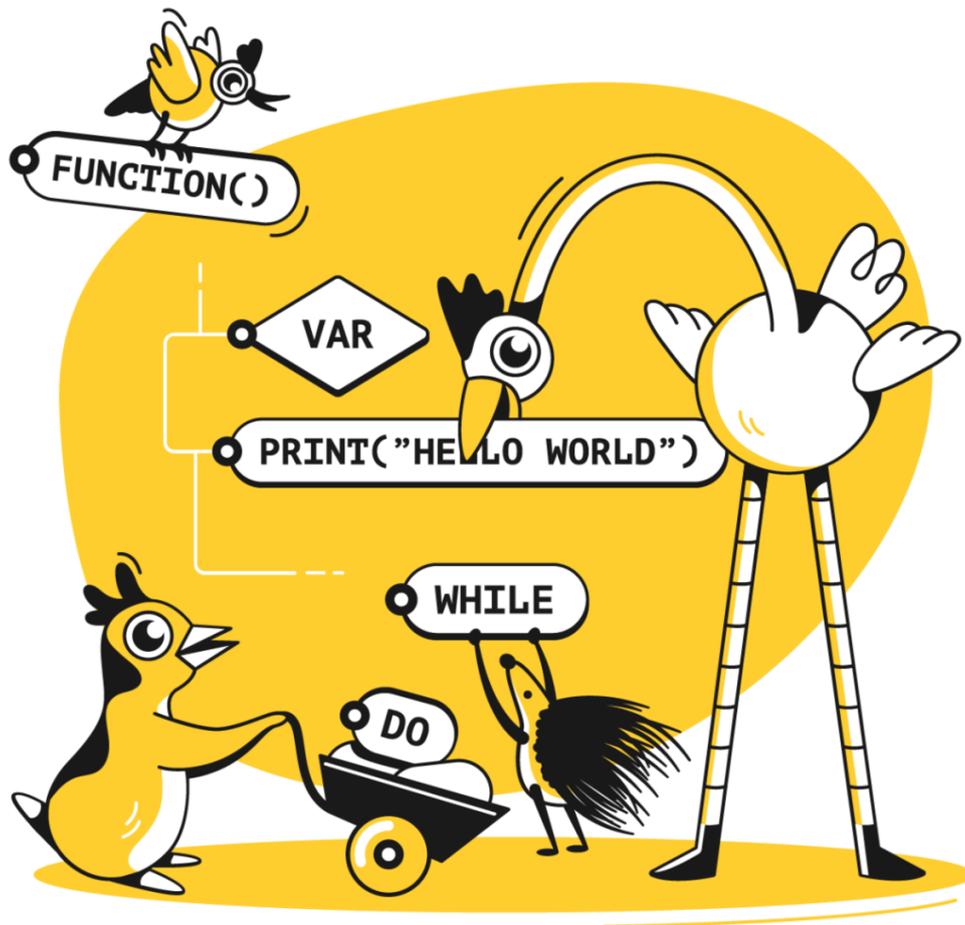
So, You will have to write that system and make sure that it works flawlessly.

In other words, You have to understand the whole ecosystem — preferences, menu, pricing, purchase orders, transactions details, payment gateway, mode of payment, CRM, time stamps, etc.

Now, let's dig a little deeper to understand the real details of what it takes to be a Machine Learning Scientist/ Engineer.

— Computer Science Fundamentals

Computer science fundamentals are exceedingly important for Data Scientists, Machine Learning Scientists and Engineers. You need to have a solid understanding of [data structures](#), [algorithms](#), [computability](#) and complexity as well as [computer architecture](#).



So, You must have a good knowledge and Skills in [Run-time Analysis of algorithms](#), [Hash Table](#), [Trees](#) (Data Structures), [Object-Oriented Programming](#) (OOP)

Practical Learning Resources

Here are some excellent resources to hone your skills.

[Fundamentals of Computing](#) (Highly recommended)

[Accelerated Computer Science Fundamentals](#) (Intermediate learners)

— Probability and Statistics for Machine Learning

Machine learning is all about predictions, supervised learning, unsupervised learning, etc. While, Statistics is about sample, population, hypothesis, etc.

[Robert Tibshirani](#), Professor in the Department of Statistics and Biomedical Data Science at Stanford University., calls machine learning “[glorified statistics.](#)”



Machine learning and statistics have the same objective.

Machine learning is a subfield of computer science and artificial intelligence. It

deals with building systems that can learn from data, instead of explicitly programmed instructions. Whereas, A statistical model, on the other hand, is a subfield of mathematics.

[Normal Deviate](#), In his blog post “[Machine Learning vs Statistics](#)” states how the same concepts have different names in the two fields

Statistics	Machine Learning
Estimation	Learning
Classifier	Hypothesis
Regression	Supervised Learning
Classification	Supervised Learning
Data Point	Example/ Instance
Covariate	Feature
Response	Label

In short, Machine Learning algorithms are essentially extensions of statistical modelling procedures. So, Machine Learning practitioners need to have a solid understanding of statistics.

Practical Learning Resources

Here are some excellent resources to intellectually bootstrap your understanding of Statistics required in Machine Learning.

Probability and [Statistics for Data Science](#) (beginners)

[Statistical Thinking in Python](#)

[Preparing for Statistics Interview Questions in Python](#)

[Statistical Modelling in R](#)

[Spatial Statistics in R](#)

[Interactive Classes, Interview Questions and Projects](#) by DataCamp

— Data Modeling, Metrics and Evaluation

Data modeling is the process of estimating the underlying structure of a given dataset, with the goal of finding useful patterns and requirements needed to support the business processes.

Your skills will apply in [finding correlations](#), [clusters](#), [eigenvectors](#), etc. and/or predicting properties of previously unseen instances such as [classification](#), [regression](#), [anomaly detection](#), etc

As a Machine Learning Engineer your success and intellectual growth do not depend on creating machine learning models. In fact, You will create Machine learning models through out your career but your growth is primarily contingent on your rigorous approach towards Metrics for evaluation, Estimation and even how you evaluate those models.



Depending on the task at hand, you will need to be able to make choices whether to select an appropriate accuracy or error measure. So, You must also have an advanced familiarity with the common metrics for evaluating models, completely depending on whether you are working with a regressing model or a classification model.

Mean-squared error

R-Squared

Accuracy

ROC AUC

Log Loss

Some of these metrics are often used in competitions like those run by [Kaggle](#).

Furthermore, A key part of the estimation process is continually evaluating how good a given model is and evaluation strategy is an essential process that includes training-testing split, sequential vs. randomized cross-validation, etc.

Also, Iterative learning algorithms often directly utilize resulting errors to tweak the model (e.g. back-propagation for neural networks), so understanding these measures is very important even for just applying standard algorithms.

Practical Learning Resources

Mentor Led [Machine Learning Career Track](#) (Job Guarantee)

[Introduction to Machine Learning](#) with R (DataCamp)

[Machine Learning Course](#) by Andrew NG

— Applied Machine Learning: Algorithms and Libraries

Machine Learning is all about Algorithms and Machine Learning Engineers are well versed with programming them into computers.

In Machine Learning, more often, our goal is prediction. For example, We can predict the sale price of a house by using a set of characteristics related to that house.

This process is simply accomplished by estimating the value of an output variable from a set of input variables.



Skilled researchers are also well versed with the most prominent algorithms that come in three groups: linear-models, tree-band models, and neural networks.

You can take a look at these major types of [machine learning algorithms in Data Science Skills](#) article I wrote last month.

Now, let's quickly dive into the important, implementation part of Machine Learning algorithms.

Today, the algorithms we need are widely available through libraries/packages/APIs ([scikit-learn](#), [TensorFlow](#), [Spark MLlib](#) and [many more](#)) but applying them effectively require skills.

You can find great resources on [Kaggle](#) to get exposed to applying [machine learning algorithms](#).

Thanks to Open-Source community and Machine learning enthusiasts who sacrificially devote their time to make these libraries available for anyone.

Practical Learning Resources

[Machine Learning BootCamp](#) (Unlimited 1:1 Coaching)

[Mathematics for Machine Learning](#) (Beginners)

[Algorithms, Data Collection and Starting to Code](#) (Specialization)

[Learn Algorithms](#) (Interactive Courses)

— Feature Engineering: Software and System Design

All machine learning algorithms use some inputs data to create output.

Machine Learning engineer's typical output or deliverable is Software and System Design.

From my interviews with Mentors at SpringBoard for [Machine Learning Career Track](#), they emphasised on the main goals of feature engineering efforts that they help their learners understand;

How to prepare the proper input dataset, compatible with the machine learning algorithms requirements to effectively work with the system design.

Learning the validation process to carefully examine each steps in Improving the performance of machine learning models through the rigorous process of validation, where you optimize the parameters for each algorithm you want to use.



If you want to get ahead in your Machine Learning career, You need to be equipped with sounds skills to understand; library calls, REST APIs, database queries, etc. and have functional skills to build appropriate interfaces for your component that others will depend on.

Rigorous system design is necessary to avoid bottlenecks and let your algorithms scale well with increasing volumes of data to run on lower computational costs.

Lastly, you need to tighten your grips on the best practices of software engineering including but not limited to [requirements analysis](#), [system design](#), [modularity](#), [efficient version control](#), [testing](#), [documentation](#), etc. These skills are exceedingly helpful for agility, productivity, collaboration, quality and maintainability.

Practical Learning Resources

Learn [Feature Engineering](#) (Google Cloud)

[TensorFlow Courses from Notable Educators](#) (Intermediate Learners)

Thanks for making it to the end ;)

If you liked this article, I've got a few recommendations for you. One about the [Machine Learning Career Track](#) by Springboard and one about [Best Machine Learning Courses](#) (from World-Class Educators) to hone your craft.

I've also got this [data science and ai newsletter](#) that you might be into. I send a tiny email once every fortnight with some useful and cool stuff I've found/made.

Don't worry, I hate spam as much as you. 🙏

Citations

Chakraborty, A. (2016, April 7). 5 Skills You Need to Become a Machine Learning Engineer. Retrieved from <https://blog.udacity.com/2016/04/5-skills-you-need-to-become-a-machine-learning-engineer.html>

Shah, Aatash. "Machine Learning vs Statistics." *kdnuggets.com*.
<https://www.kdnuggets.com/2016/11/machine-learning-vs-statistics.html> (accessed November 21, 2019).

Kelley, R. (2017, January 19). *Machine Learning Explained: Algorithms Are Your Friend*. Retrieved from <https://blog.dataiku.com/machine-learning-explained-algorithms-are-your-friend>

Image Credits:

eSparbiz

[/dribbble.com/shots/3517348-Not-Hotdog](https://dribbble.com/shots/3517348-Not-Hotdog)

qz.com/1261817/predictive-algorithms-are-not-all-that-complicated/

dribbble.com/crispe

dribbble.com/Ivan_Haidutski

dribbble.com/kajdax