

Before we get started:

The image illustrates a workflow for opening a file in Google Drive. It is divided into three main sections:

- Mobile Interface:** A screenshot of a mobile app showing an "In-call messages" dialog. A message from "You" at 12:30 PM contains a link: https://drive.google.com/file/d/1Xqac_Y8vDpAFzA_KI_SOLDDBZxs9F5YHh/view?usp=drive_link. A hand icon points to the link, and a thick black arrow points upwards to the next section.
- Desktop File Menu:** A screenshot of a "File" menu. The menu items are: "Locate in Drive", "Open in playground mode", "New notebook", "Open notebook" (with "Ctrl+O" shortcut), "Upload notebook", "Rename", "Move", "Move to trash", "Save a copy in Drive" (highlighted with a black box and a hand icon), "Save a copy as a GitHub Gist", and "Save a copy in GitHub". A thick black arrow points from the mobile interface to this menu.
- Text:** The text "om SC" is written in a large, bold, black font. Below it, the text "From A to Z. From zero t" is partially visible.


NLP from scratch



Machine Learning: From Zero

April 2nd, 2025

www.nlpfromscratch.com



Housekeeping



Camera on if comfortable doing so



This meeting will not be recorded



Stay muted unless speaking



Be professional

Who am I?

- Data Scientist
- Career consultant (SapientNitro, PwC, Accenture)
- Community builder
- Accidental entrepreneur




Manifesto



Knowledge is only valuable if it is useful.



The best way to learn is by doing.



Learning is a non-linear process.




Learning is exploration, not a journey.



Teaching and learning are complementary.

Tools of the Trade



Python 3

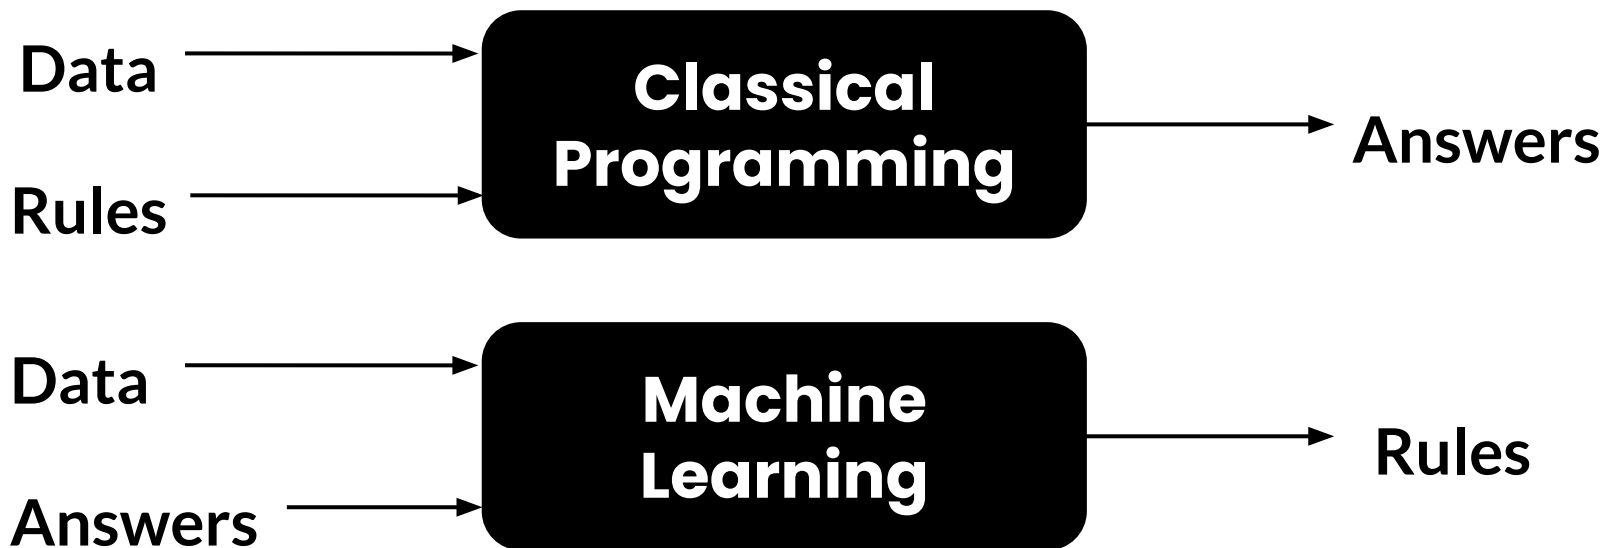


Google Colab
/ Jupyter




Machine Learning

What is Machine Learning?




Types of Machine Learning



Supervised Learning

Making predictions

Make predictions from a dataset and *data labels* - categorical or numeric values associated with each observation in the training data.



Unsupervised Learning

Finding patterns in data

Uses statistical techniques to uncover patterns in a dataset based upon its features. Unlike supervised learning, it does not require data labels.



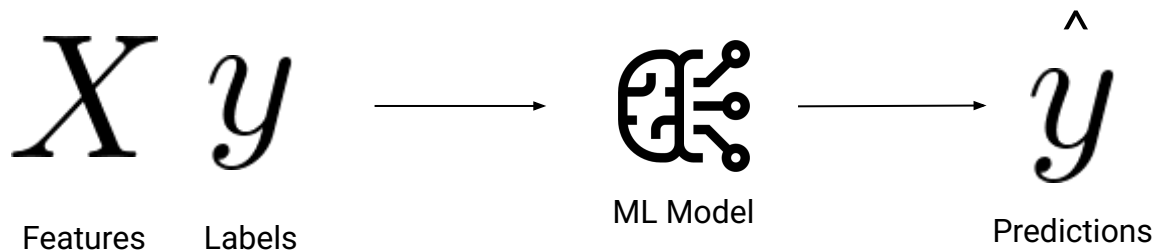
Reinforcement Learning

Feedback & Rewards

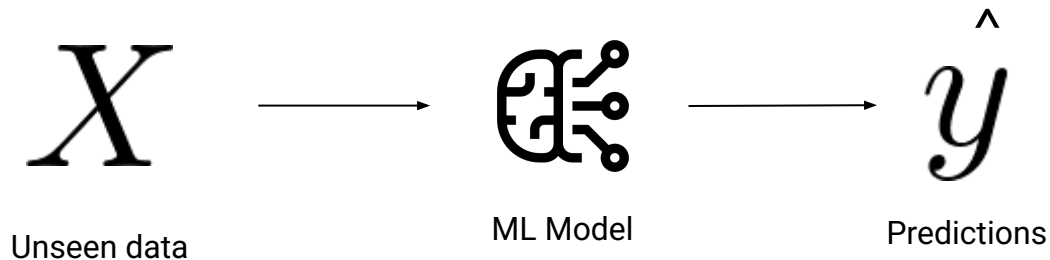
Reinforcement Learning (RL) teaches an *agent* a behavior by optimizing against a target objective with a reward function.

Supervised Learning

TRAINING



PREDICTION / EVALUATION




$$\text{Error} = y - \hat{y}$$

Linear Models

Regression

Linear Regression
(Ordinary Least Squares / OLS)

$$\hat{y} = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \dots$$




Classification

Logistic Regression

$$Y = \ln\left(\frac{p}{1-p}\right) = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_n X_n$$

$$\ln\left(\frac{p}{1-p}\right) = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_n X_n$$

$$P = \frac{e^{a+bX}}{1 + e^{a+bX}}$$




A Machine Learning Problem



Overfitting


The goal of a given ML model is to be as generalizable as possible and describe the underlying behavior or phenomenon of interest.

Overfitting refers to the model learning specifics of the particular dataset that was used for training, such that it does not perform well on new data.




Overfitting:

The model has learned "too much" about the dataset used for training, including outlying observations and is fit too tightly to the data.



Good Fit:

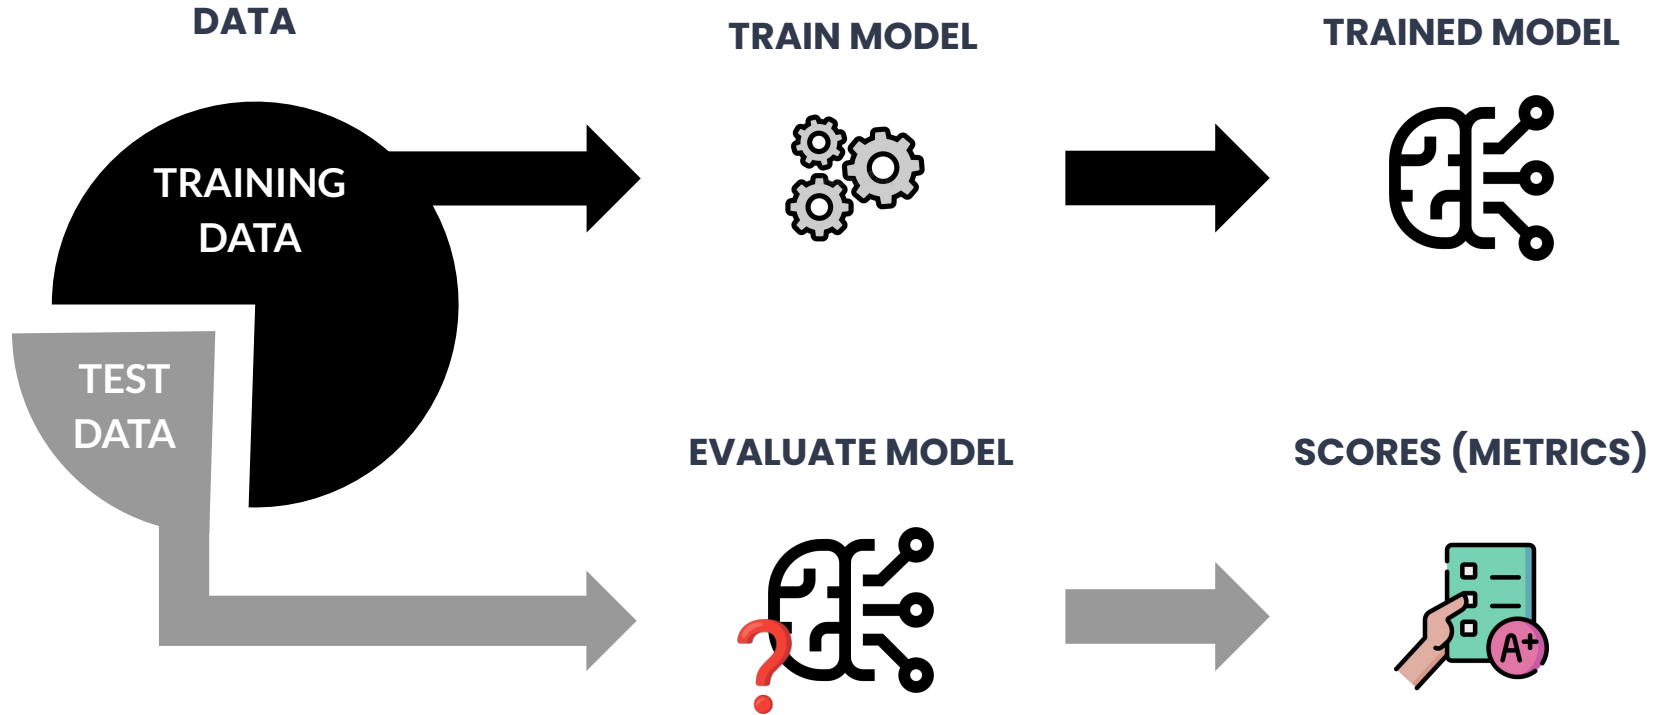
The model is fit well, describing the underlying pattern in the data well in a general way. Its predictions are not perfect, but it does not exhibit excessive variance in order to capture all data points.




Underfitting:

The model has not learned enough about the patterns in the data and fails to adequately capture the relationship it is supposed to describe.

Training and Testing an ML Model from Data



NLP from scratch 

QUESTIONS



NLPfromscratch.com

🔍 Search NLP from scratch



NLP from scratch is a collection of free and Pay-What-You-Can (PWYC) courses and workshops created by [Myles Harrison](#).

This official homepage contains all the files and relevant resources, hosted on Github.

- [Webinars](#): Webinars on NLP, LLMs, and OpenAI delivered from October 2023 - Nov 2024.
- [Workshops](#): Various workshops on Generative AI and related topics.
- [Datasets](#): NLP / language datasets used in the above.
- [Talks](#): Recordings of talks at various meetup events and conferences.

If you find the materials I've created useful, please consider supporting me in developing more by contributing a donation:

Donate

NLP from scratch



www.nlpfromscratch.com