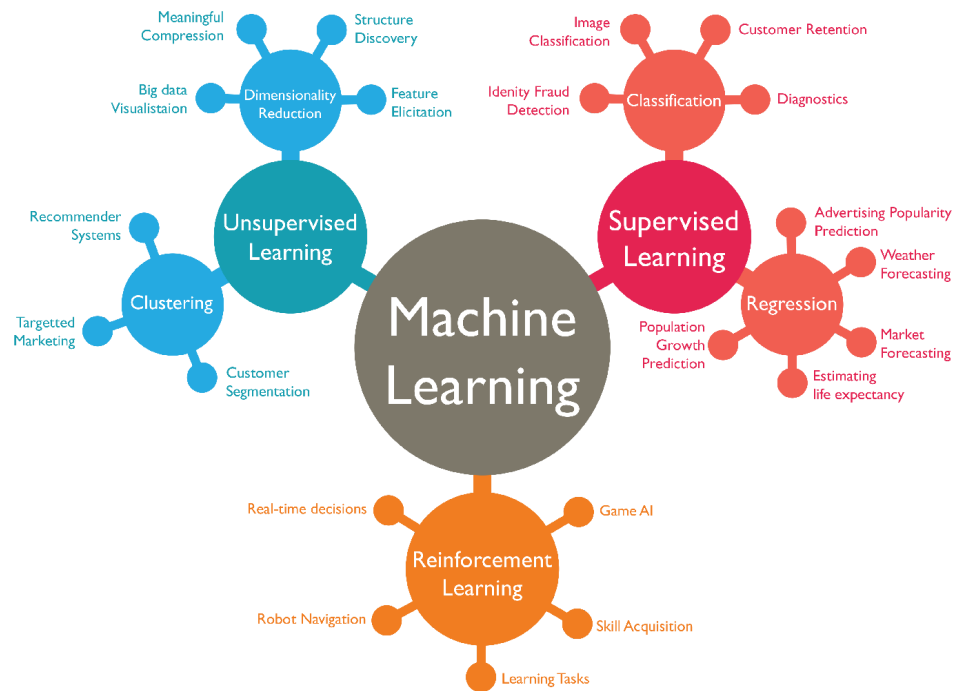
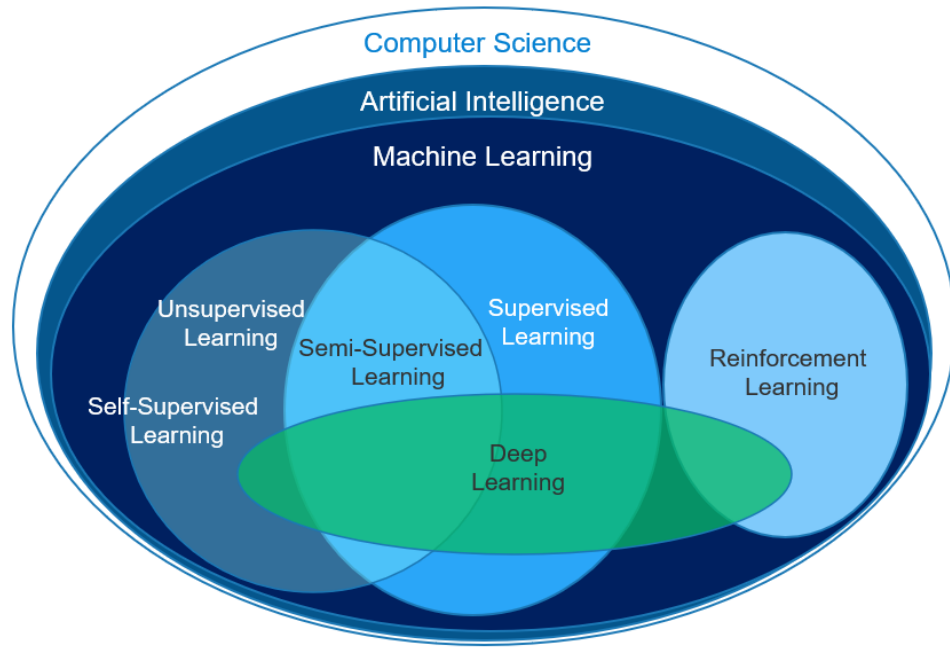


Machine Learning Overview



Maching Learning vs Deep Learning

Machine Learning vs. Deep Learning

(common algorithms)

- Random forest
- Naive bayes
- Nearest neighbour
- Support vector machine
- ...many more

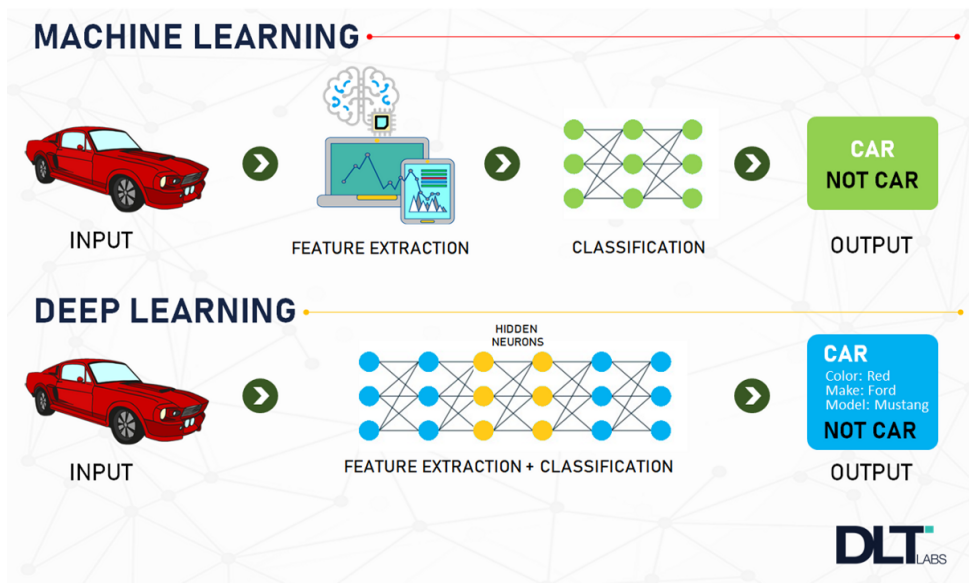
(since the advent of deep learning these are often referred to as "shallow algorithms")

- Neural networks
- Fully connected neural network
- Convolutional neural network
- Recurrent neural network
- Transformer
- ...many more

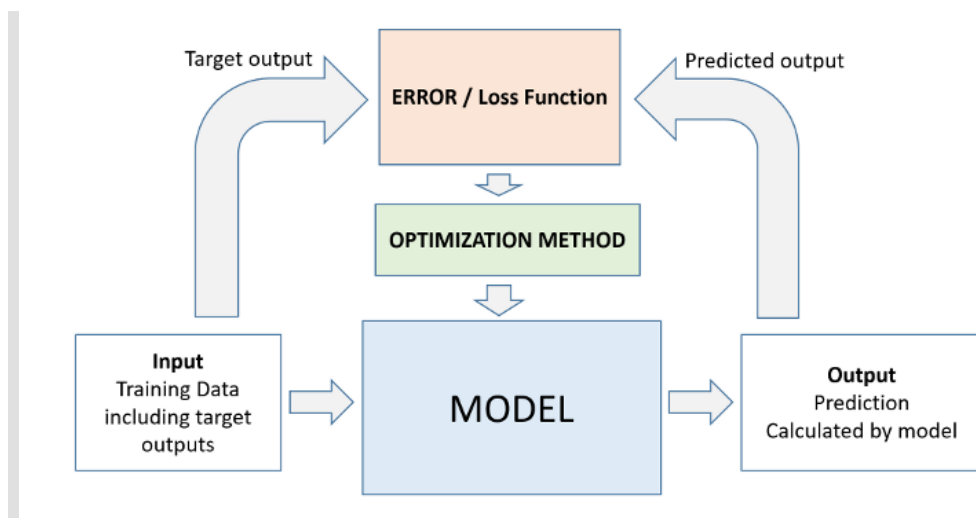
What we're focused on building (with TensorFlow)

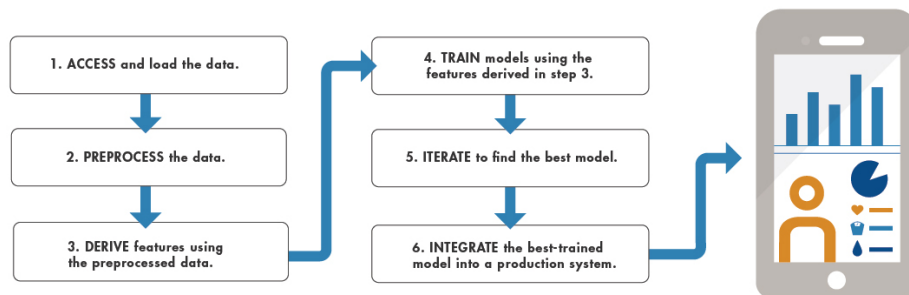
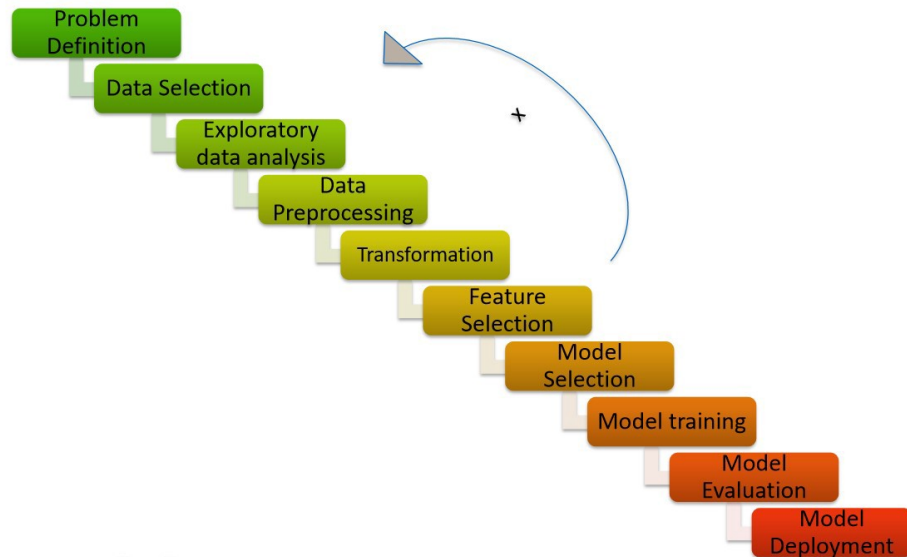
(depending how you represent your problem, many algorithms can be used for both)

Structured data ← → Unstructured data



Machine Learning Workflow





Workflow 1. Access and load Data

```
import opendatasets as od
```

```
-> od.download(dataset_url)
```

```
from urllib.request import urlretrieve
```

```
-> urlretrieve(url, "path/file") # rewrite the file (becareful)
```

```
from zipfile import ZipFile
```

```
-> with ZipFile('filename.zip') as f:
```

```
f.extractall(path='filename')
```

```
import pandas as pd
```

```
-> pd.read_csv("path/csvfile")
```

Workflow 2. Preprocess the Data -> Design Thinking

2.1 Exploratory Analysis

```
import numpy as np
```

```
-> np.arange()
```

-> `np.mean()`

```
import pandas as pd
```

-> `df.head()`,

-> `df.sample()*`

-> `df.isna().sum()`

-> `df.unique()`

-> `df.nunique()`

-> `df.value_counts()`


-> `df.sum()`

-> `df.shape`


-> `df.info()`

-> `df.describe()`


-> `df.corr()` ([Source](#)) > ([Correlation Explanation](#))




Perfect
Positive
Correlation




Strong
Positive
Correlation




Weak
Positive
Correlation




No
Correlation



Weak
Negative
Correlation



Strong
Negative
Correlation



Perfect
Negative
Correlation

2.2 Visualization

```
import matplotlib.pyplot as plt
```

```
import seaborn as sn
```

```
import plotly.express as px
```

2.3 Imputing Missing Numeric Data

- `pandas -> pd.DataFrame -> df`
 -> `df.to_datetime()`
 -> `df.fillna()`
 -> `df.dropna(subset=[])`
- `from sklearn.impute import SimpleImputer`

| | col1 | col2 | col3 | col4 | col5 | | col1 | col2 | col3 | col4 | col5 | |
|---|------|------|------|------|------|----------|------|------|------|------|------|-----|
| 0 | 2 | 5.0 | 3.0 | 6 | NaN | → mean() | 0 | 2.0 | 5.0 | 3.0 | 6.0 | 7.0 |
| 1 | 9 | NaN | 9.0 | 0 | 7.0 | | 1 | 9.0 | 11.0 | 9.0 | 0.0 | 7.0 |
| 2 | 19 | 17.0 | NaN | 9 | NaN | | 2 | 19.0 | 17.0 | 6.0 | 9.0 | 7.0 |

Workflow 3. Drive Features Using the Processed Data

- `Numerical features` : StandardScaler or MinMaxScaler
- `Categorical features` OneHotEncoder
- `Using Power BI to Prepare the Data`

Numerical features

- StandardScaler

```
from sklearn.preprocessing import StandardScaler
-> df=StandardScaler().fit_transform(df)
```

Standardization:

$$z = \frac{x - \mu}{\sigma}$$

with mean:

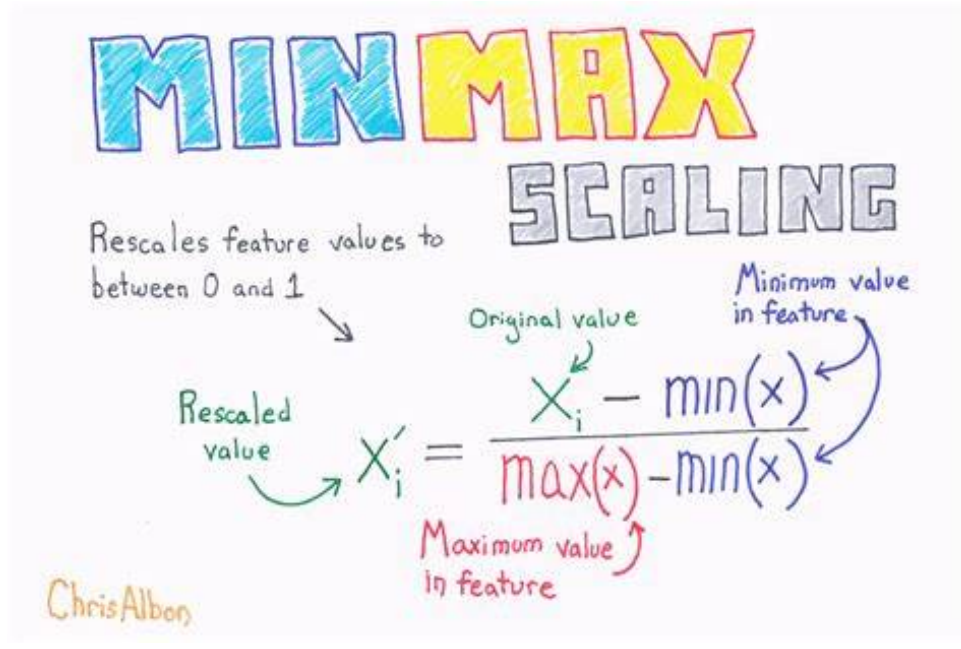
$$\mu = \frac{1}{N} \sum_{i=1}^N (x_i)$$

and standard deviation:

$$\sigma = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - \mu)^2}$$

- MinMaxScaler

```
from sklearn.preprocessing import MinMaxScaler
-> df=MinMaxScaler().fit_transform(df)
```



Categorical Features

```
df.map({})
```

```
from sklearn.preprocessing import OneHotEncoder
```

One hot encoding involves adding a new binary (0/1) column for each unique category of a categorical column.

| Index | Categorical column |
|-------|--------------------|
| 1 | Cat A |
| 2 | Cat B |
| 3 | Cat C |

→

| Index | Cat A | Cat B | Cat C |
|-------|-------|-------|-------|
| 1 | 1 | 0 | 0 |
| 2 | 0 | 1 | 0 |
| 3 | 0 | 0 | 1 |

```
pd.get_dummies(df['columnName'])
```

Pandas Get Dummies

Turn your Categorical Column (Ex: "Name")...

| Index | Name | 8/6/2020 |
|-------|------------|----------|
| 0 | Liho Liho | \$234.54 |
| 1 | Chambers | \$45.74 |
| 2 | The Square | \$56.22 |
| 3 | Liho Liho | \$32.31 |

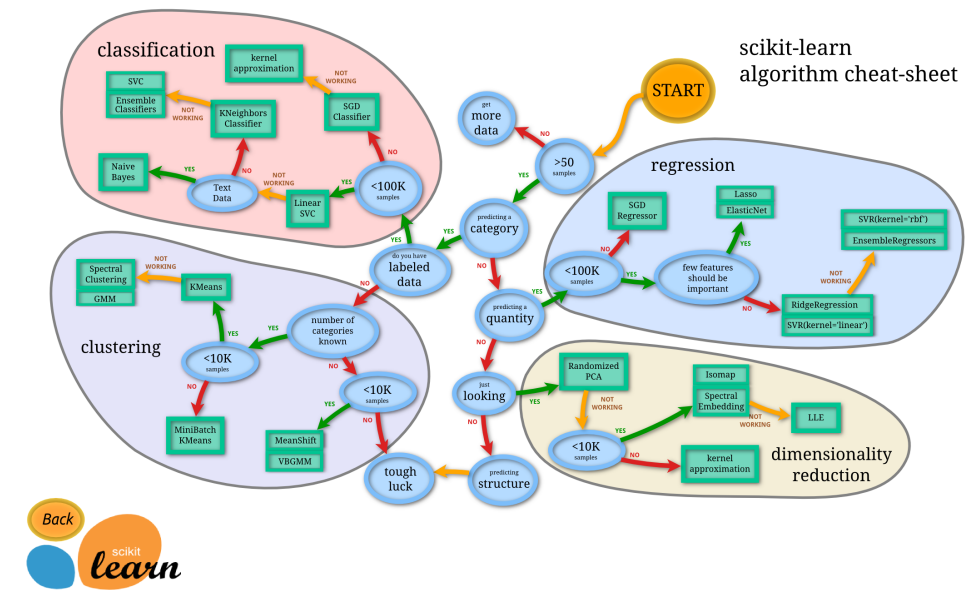
...Into Dummy Indicator Columns

| Index | Liho Liho | Chambers | The Square | 8/6/2020 |
|-------|-----------|----------|------------|----------|
| 0 | 1 | 0 | 0 | \$234.54 |
| 1 | 0 | 1 | 0 | \$45.74 |
| 2 | 0 | 0 | 1 | \$56.22 |
| 3 | 1 | 0 | 0 | \$32.31 |

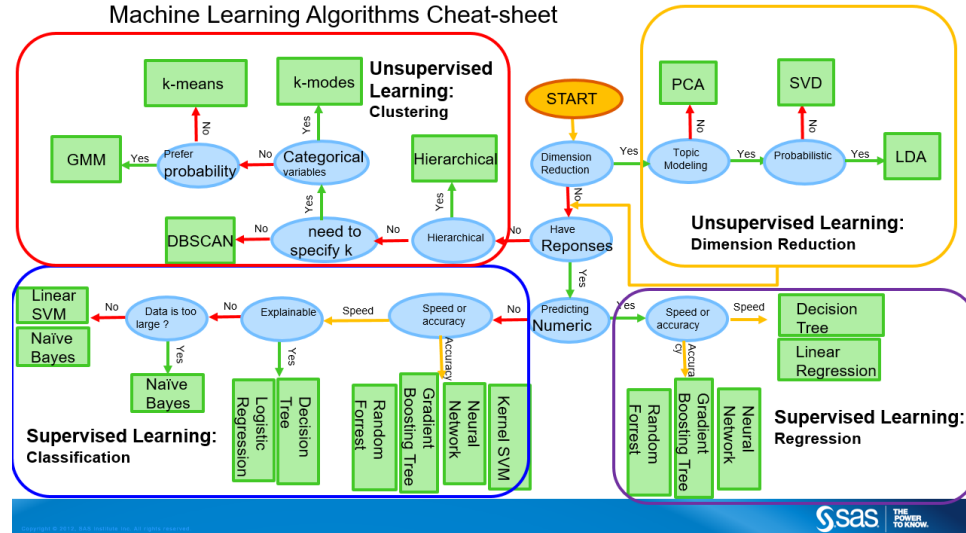
Workflow 4. Model Selection and Training

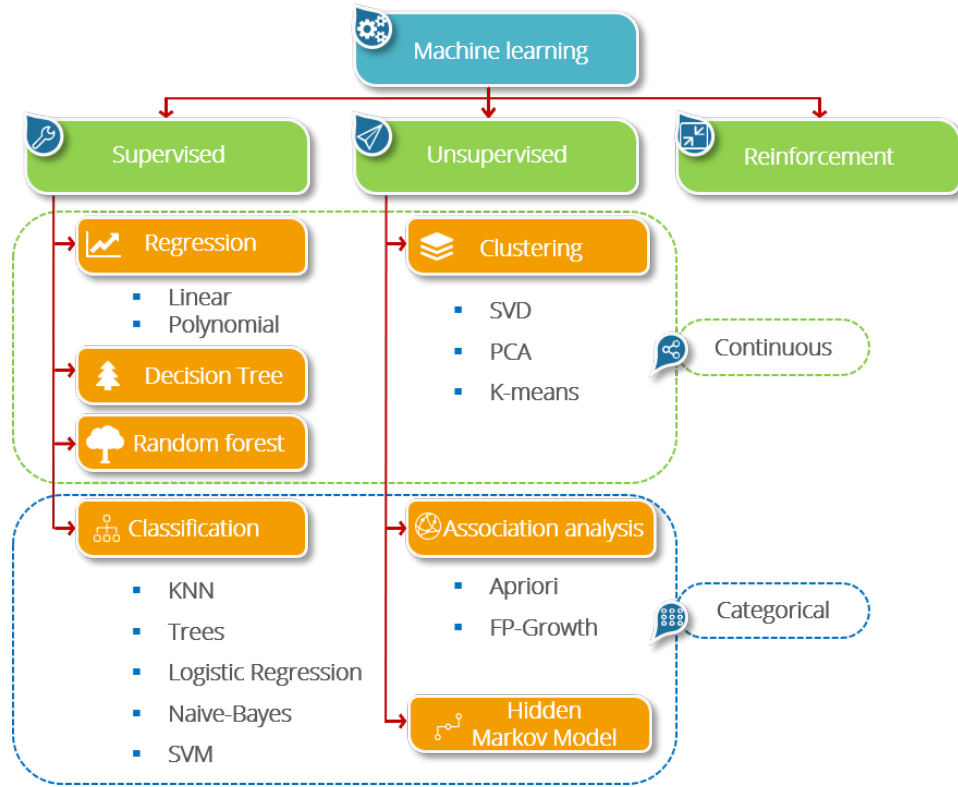
4.1 Model Selection

Scikit-learn offers the following cheatsheet to decide which model to pick for a given problem



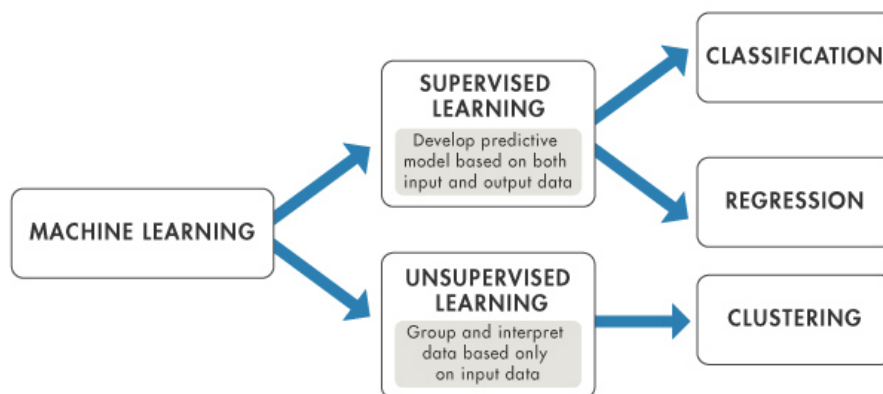
Machine Learning Algorithms Cheat-sheet



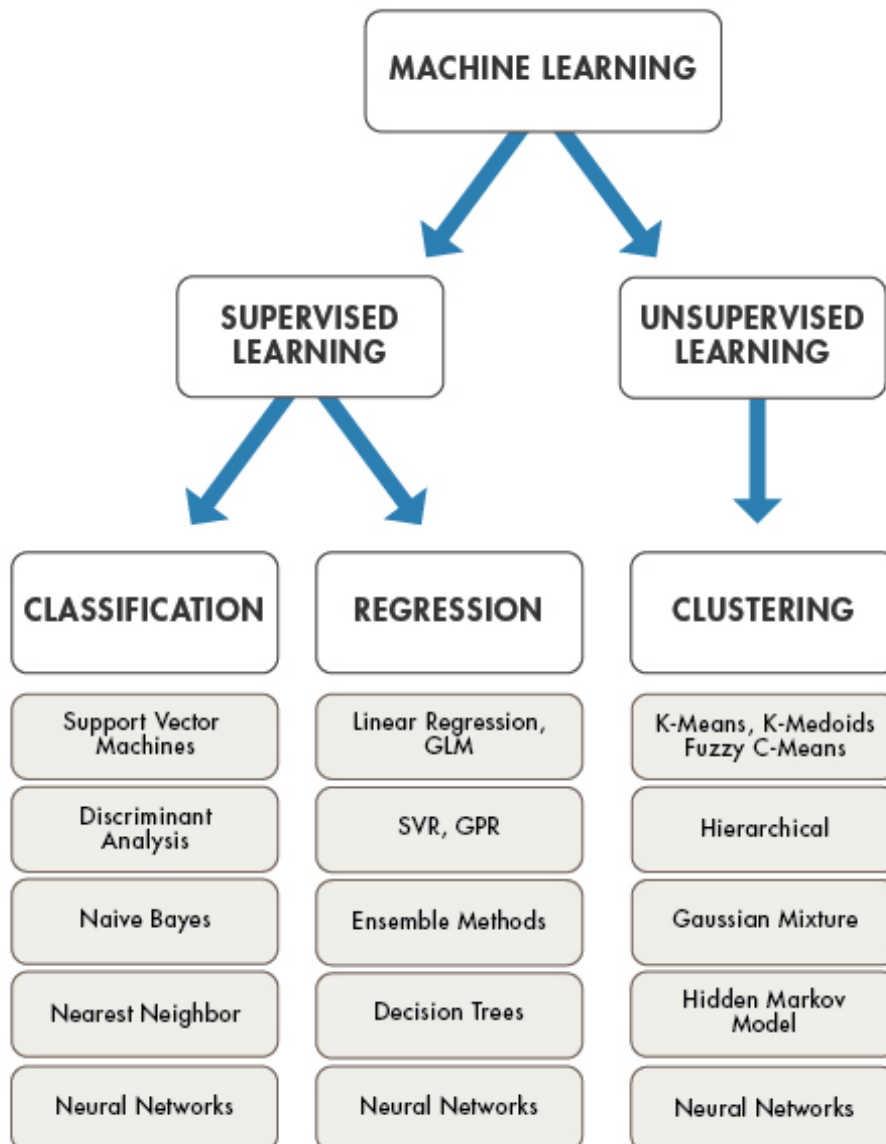
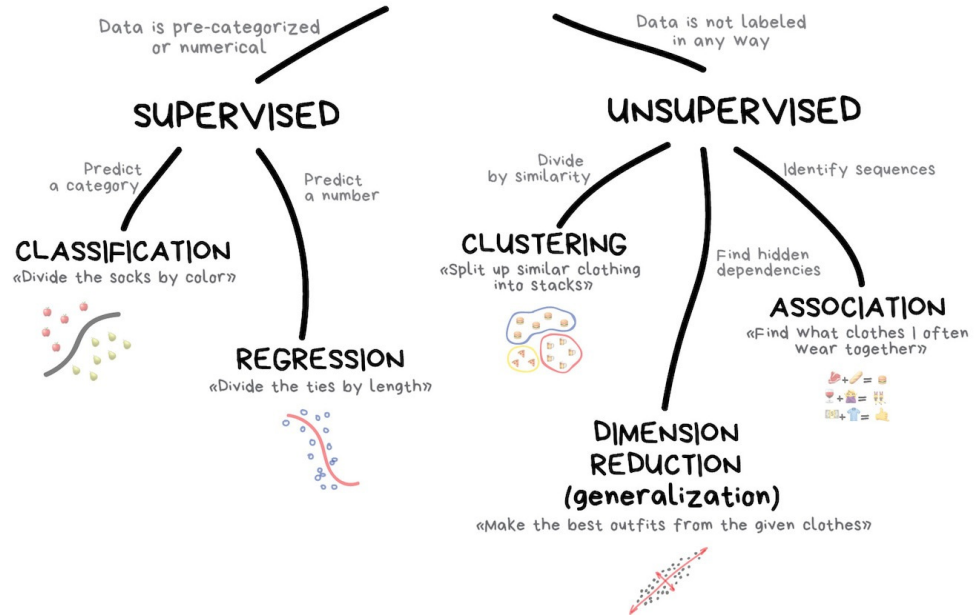


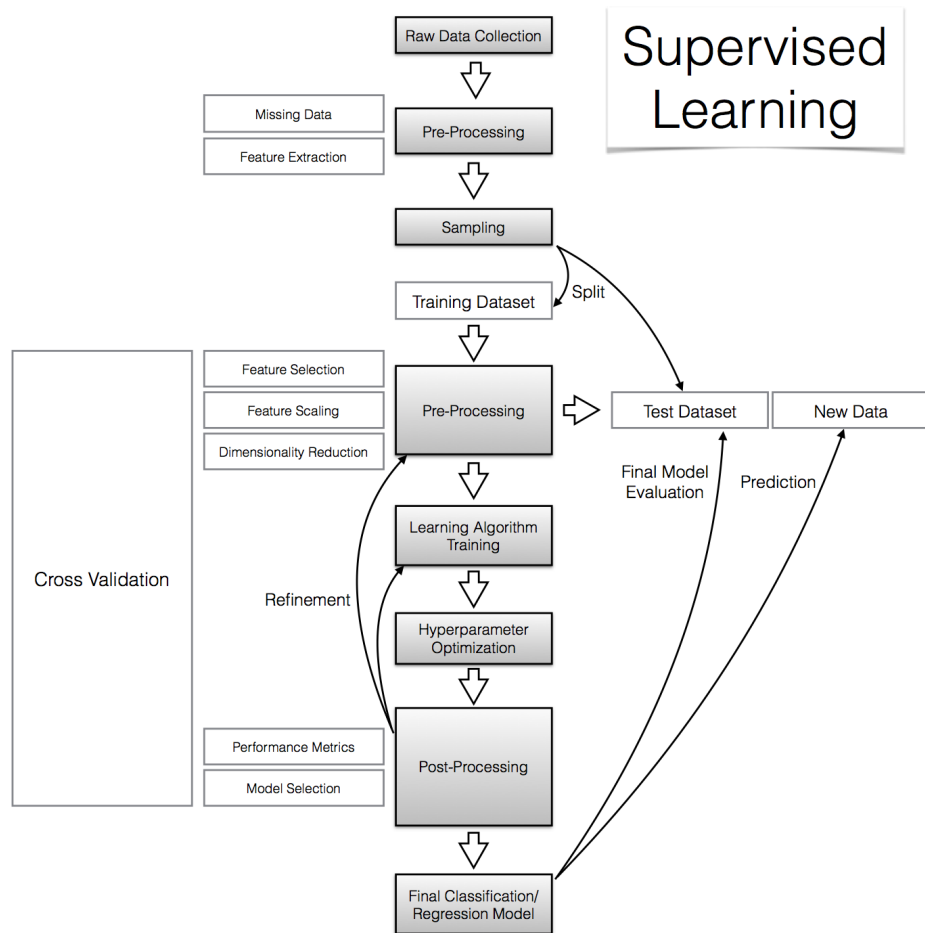
Supervised learning vs Unsupervised learning

Machine learning uses two types of techniques:(image source) -> Supervised learning, which trains a model on known input and output data so that it can predict future outputs. -> Unsupervised learning, which finds hidden patterns or intrinsic structures in input data.



CLASSICAL MACHINE LEARNING





Supervised Learning -Regression Problems


```
from sklearn.linear_model import LinearRegression
from sklearn.linear_model import SGDRegressor
from sklearn.linear_model import Ridge
```

```
-> *model = LinearRegression().fit(inputs, targets)*
-> *model = SGDRegressor().fit(inputs, targets)*
-> *model = Ridge().fit(inputs, targets)*
-> *model.coef_, model.intercept_*
```

$$\rightarrow y = w \times x + b$$

- The numbers w and b are called the **parameters** or **weights** of the model.

Here's a visualization of how gradient descent works:

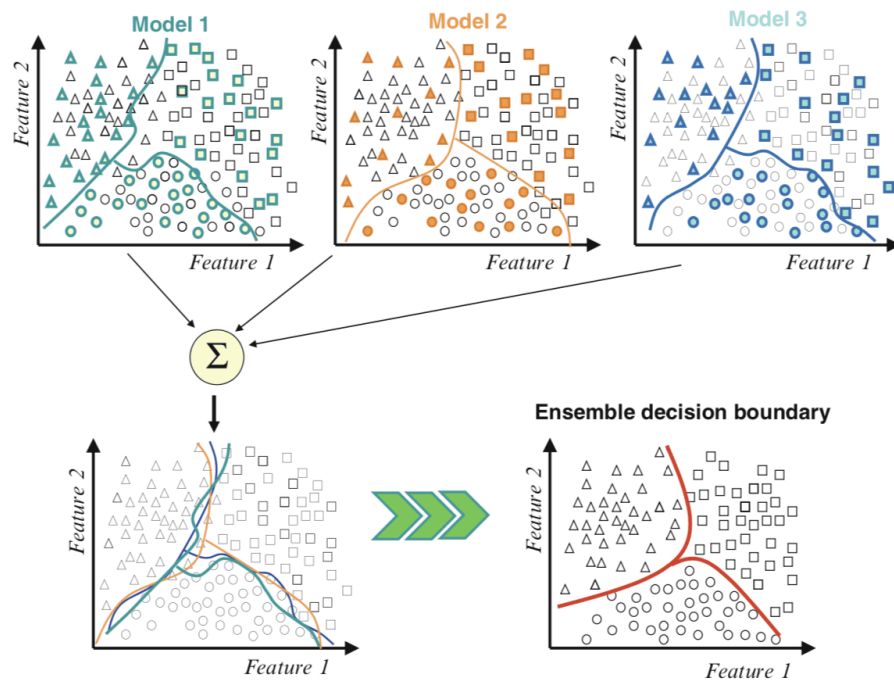


Random Forest

```
from sklearn.ensemble import RandomForestClassifier
```

```




```



Supervised Learning - Classification Problems

Logistic Regression

```
from sklearn.linear_model import LogisticRegression
```



Decision Tree Classifier

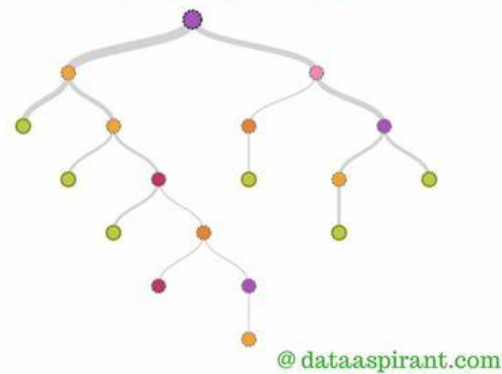
```

from sklearn.tree import DecisionTreeClassifier
-> Feature Important
-> np.argmax(model.feature_importances_)
-> df_fi = pd.DataFrame({'Feature': inputs.columns,
'Importance':
model.feature_importances_}).sort_values('Importance',
ascending=False)
from sklearn.tree import plot_tree
from sklearn.tree import export_text

```


Building Decision Tree Classifier

In Python with Scikit-Learn




Support Vector Machine

```
from sklearn.svm import SVC
```



KNN K nearest neighbors Classification


```
from sklearn.neighbors import KNeighborsClassifier -> knn =  
KNeighborsClassifier(n_neighbors=10)
```



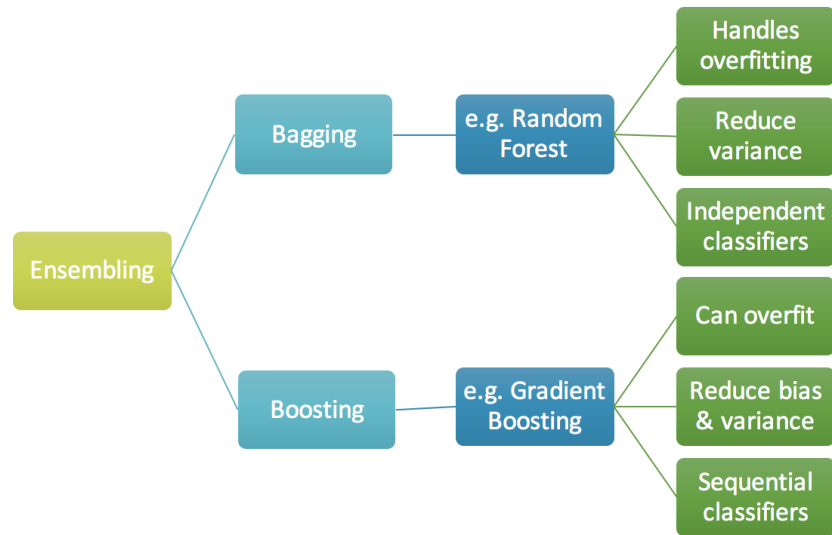
Naive Bayes Classifier

$$P(H|E) = \frac{P(E|H) * P(H)}{P(E)}$$

Likelihood of the Evidence given that the Hypothesis is True (points to $P(E|H)$)
Prior Probability of the Hypothesis (points to $P(H)$)
Posterior Probability of the Hypothesis given that the Evidence is True (points to $P(H|E)$)
Prior Probability that the evidence is True (points to $P(E)$)



Gradient Boosting Classifier




Gradient Boosting (Simple Version)


(Why is it called “gradient”?)

(For Regression Only)

(Answer next slides.)




<http://statweb.stanford.edu/~jhf/ftp/trebst.pdf>




Unsupervised Learning

Clustering



Dimensionality Reduction




Clustering

Clustering is the process of grouping objects from a dataset such that objects in the same group (called a cluster) are more similar (in some sense) to each other than to those in other groups ([Wikipedia](https://en.wikipedia.org/wiki/Clustering)).

Here is a full list of unsupervised learning algorithms available in Scikit-learn:


https://scikit-learn.org/stable/unsupervised_learning.html

Here is a visual representation of clustering:



DBSCAN


```
from sklearn.cluster import DBSCAN
-> model = DBSCAN(eps=0.5, min_samples=4).fit(inputs)
-> model.labels_
```



K Means Clustering

```
from sklearn.cluster import KMeans
-> model = KMeans(n_clusters=3, random_state=42).fit(inputs)
-> model.cluster_centers_
-> model.inertia_
```

The K-means algorithm attempts to classify objects into a pre-determined number of clusters by finding optimal central points (called centroids) for each cluster. Each object is classified as belonging to the cluster represented by the closest centroid.




Here's how the K-means algorithm works: 1. Pick K random objects as the initial cluster centers. 2. Classify each object into the cluster whose center is closest to the point. 3. For each cluster of classified objects, compute the centroid (mean). 4. Now reclassify each object using the centroids as cluster centers. 5. Calculate the total variance of the clusters (this is the measure of goodness). 6. Repeat steps 1 to 6 a few more times and pick the cluster centers with the lowest total variance.

Here's how the results of DBSCAN and K Means differ:



Hierarchical Clustering

Hierarchical clustering, as the name suggests, creates a hierarchy or a tree of clusters.



While there are several approaches to hierarchical clustering, the most common approach works as follows: 1. Mark each point in the dataset as a cluster. 2. Pick the two closest cluster centers without a parent and combine them into a new cluster. 3. The new cluster is the parent cluster of the two clusters, and its center is the mean of all the points in the cluster. 4. Repeat steps 2 and 3 till there's just one cluster left.


Principal Component Analysis (PCA)


```
from sklearn.decomposition import PCA
```

```
-> model = PCA(n_components=2).fit(X)
```

```
-> model.components_
```

Principal component is a dimensionality reduction technique that uses linear projections of data to reduce their dimensions, while attempting to maximize the variance of data in the projection.






t-Distributed Stochastic Neighbor Embedding (t-SNE)

```
from sklearn.manifold import TSNE
```

```
inputs_transformed = TSNE(n_components=2).fit_transform(inputs)
```

Manifold learning is an approach to non-linear dimensionality reduction. A commonly-used technique is t-Distributed Stochastic Neighbor Embedding or t-SNE.

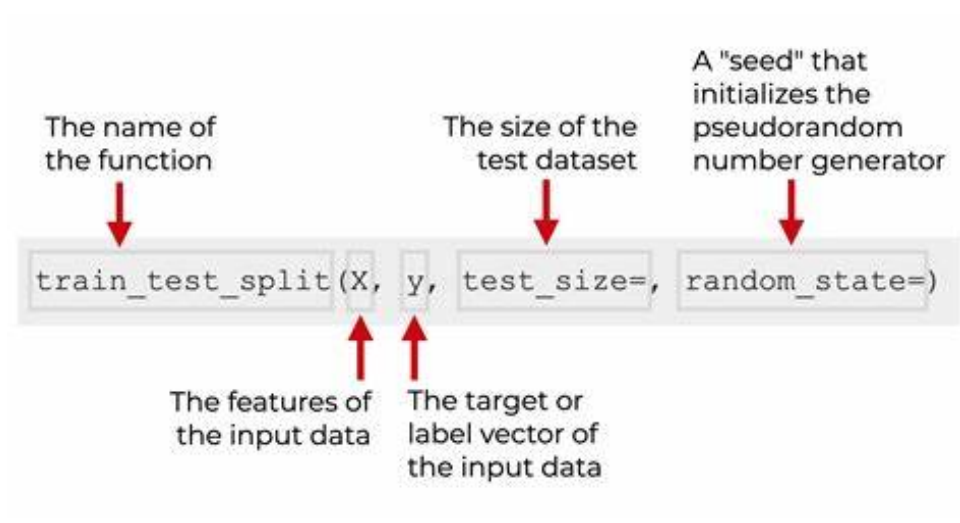
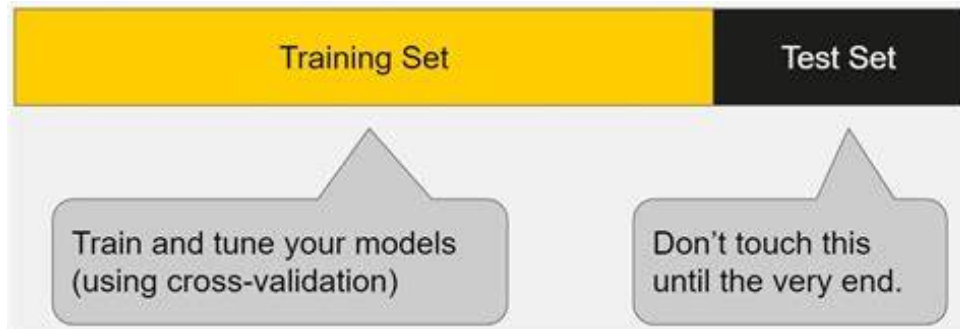


Collaborative filtering with FastAI

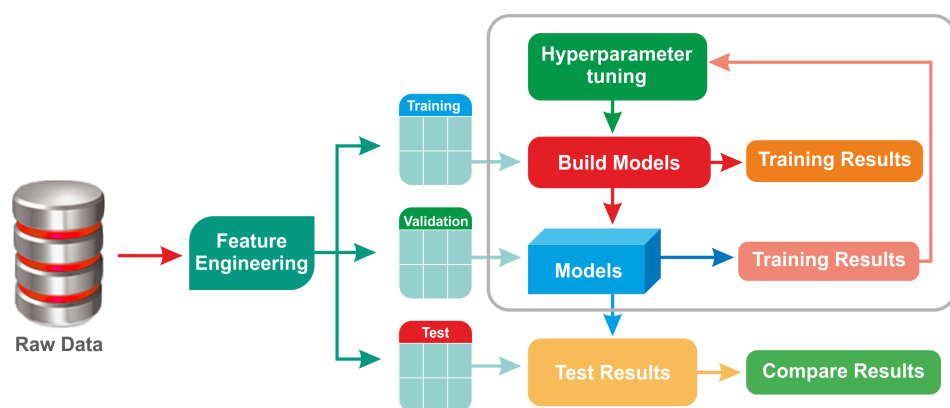
Collaborative filtering is perhaps the most common technique used by recommender systems. *Collaborative filtering is a method of making predictions about the interests of a user by collecting preferences from many users. The underlying assumption is that if a person A has the same opinion as a person B on an issue, A is more likely to have B's opinion on a different issue than that of a randomly chosen person.* - [Wikipedia](#)

4.2 Training

```
from sklearn.model_selection import train_test_split
-> X_train, X_test, y_train, y_test = train_test_split(X, y, test_size=20,
random_state=42)
```




Workflow 5. Evaluating Model



Evaluating Model - K Fold Cross Validation

```
from sklearn.model_selection import KFold
```

```
-> *for train_index, val_index in KFold().split(X):*
->     *X_train, X_val = X.iloc[train_index], X.iloc[val_index]*
->     *y_train, y_val = y.iloc[train_index], y.iloc[val_index]*
```



Evaluating Model - Score

-> `model.score(inputs, targets)`

```
from sklearn.metrics import accuracy_score
```

-> `accuracy_score(targets, predictions)`

Evaluating Model - confusion_matrix - Sensitivity and Specificity

```
from sklearn.metrics import confusion_matrix
```

Confusion Matrix



| | | <u>P R E D I C T E D</u> | |
|--------------------|------------|--------------------------|-----------------------|
| | | "Match" | "No Match" |
| <u>A C T U A L</u> | "Match" | True Positive | False Negative |
| | "No Match" | False Positive | True Negative |

| | | | |
|---------------|---------------|--|---|
| | | Predicted | |
| | | Negative (N) - | Positive (P) + |
| Actual | Negative - | True Negatives (TN) | False Positives (FP) Type I error |
| | Positive + | False Negatives (FN) Type II error | True Positives (TP) |


Evaluating Model - Loss/Cost Function

```

    stephenallwright.com


    loss = abs(actual - prediction)
    cost = sum(loss)/num_predictions
  
```

Overfitting vs. Underfitting




<img src="https://i.imgur.com/EJCrSZw.png" width=480


Bias vs. Variance



Gradient Descent



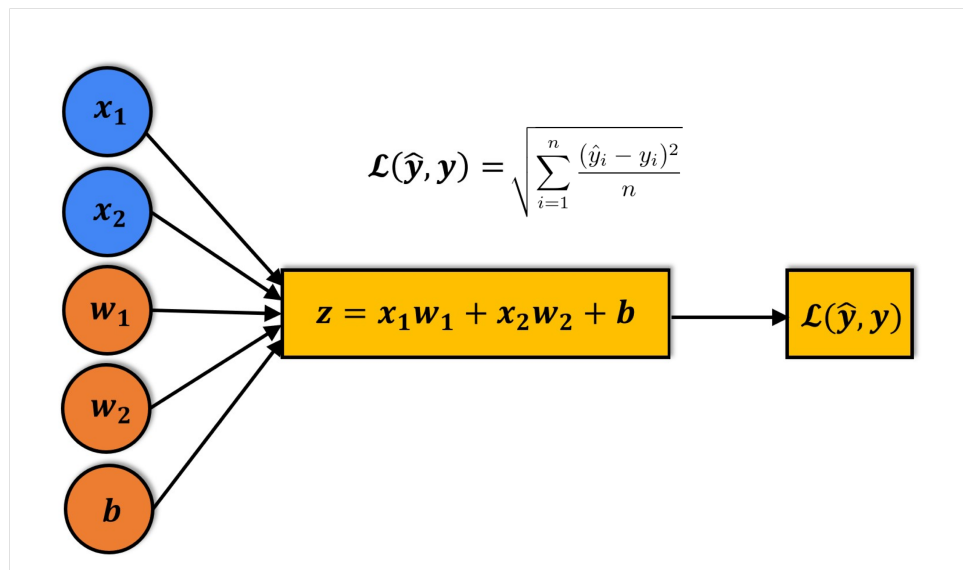
More Limitation of Gradient Descent




Regression Problems

RMSE

- > `np.sqrt(np.mean(np.square(targets - predictions)))`
- > `from sklearn.metrics import mean_squared_error`
- > `mean_squared_error(targets, predictions, squared=False)`




Geometrically, the residuals can be visualized as follows:



Classification Problems

cross entropy loss function



L1 and L2 Regularization

L1 Regularization

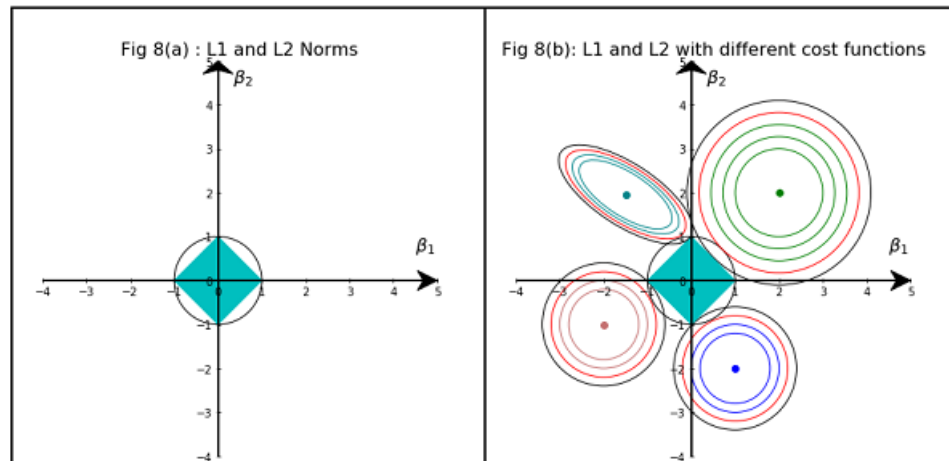
$$\text{Cost} = \sum_{i=0}^N (y_i - \sum_{j=0}^M x_{ij}W_j)^2 + \lambda \sum_{j=0}^M |W_j|$$

L2 Regularization

$$\text{Cost} = \sum_{i=0}^N (y_i - \sum_{j=0}^M x_{ij}W_j)^2 + \lambda \sum_{j=0}^M W_j^2$$

Loss function

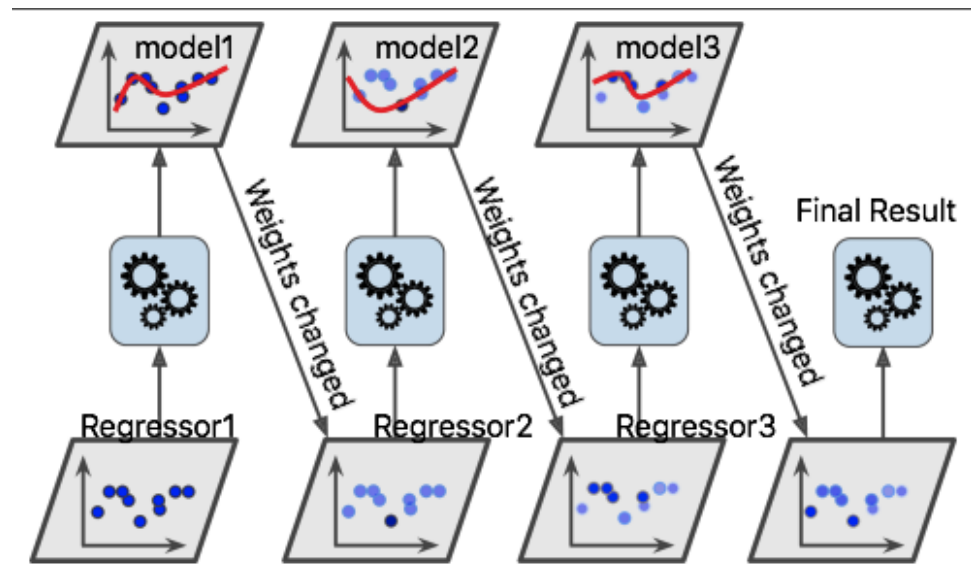
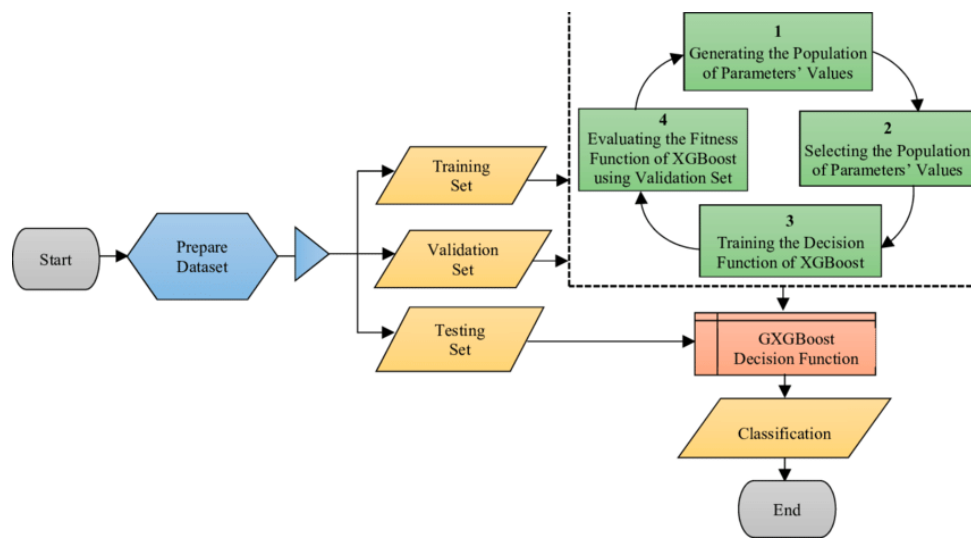
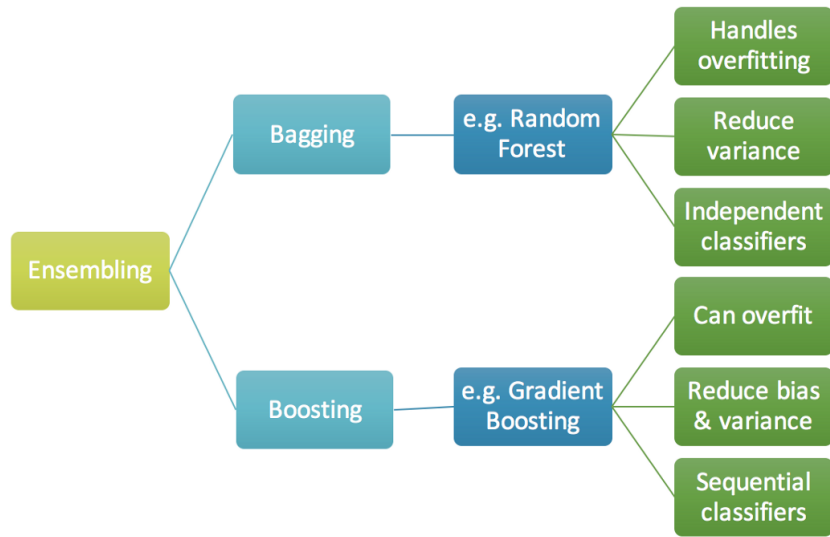
Regularization
Term




Gradient Boosting

The term "gradient" refers to the fact that each decision tree is trained with the purpose of reducing the loss from the previous iteration (similar to gradient descent). The term "boosting" refers to the general technique of training new models to improve the results of an existing model.

- [Video Tutorials on StatQuest](#)





Workflow 6. Save Model Using Joblib And Pickle

```
import pickle
```

```
-> pickle.dump(model, "path/file")
```

```
-> pickle.load("path/file")
```

```
import joblib
```

```
-> joblib.dump(model, "path/file")
```

```
-> joblib.load("path/file")
```

| Approach | PROS | CONS |
|--------------------|--|---|
| Pickle | Quick Implementation Easy Readability | Python version issues. No stored results or data Only File Object |
| Joblib | Quick Implementation Easy Readability Accepts Objects and String filenames Different compression options. | Python version issues. |
| Proprietary | More Flexibility No Python version issues | Harder Readability Complexity |

-- Memo End --