What is
ML? DL? RL? AI?

Machine Learning

- There are many domains we'll cover in this course, including:
  - ML - Machine Learning
  - DL - Deep Learning
  - RL - Reinforcement Learning
  - AI - Artificial Intelligence

Machine Learning

- This Overview Section is designed to help understand how Artificial Intelligence, Machine Learning, Deep Learning, and Reinforcement Learning are related to each other.

Machine Learning

- We'll explore "standard" machine learning concepts, such as Supervised Learning and Unsupervised Learning.
- Then we'll see how Reinforcement Learning differs from these more traditional methods.

Machine Learning

- By the end of these lectures, we'll understand the relationships between:
  - Machine Learning
  - Supervised Learning
  - Unsupervised Learning
  - Reinforcement Learning
  - Deep Learning

Machine Learning

- Let’s begin by exploring the domains:

Machine Learning

- Let’s begin by exploring the domains:

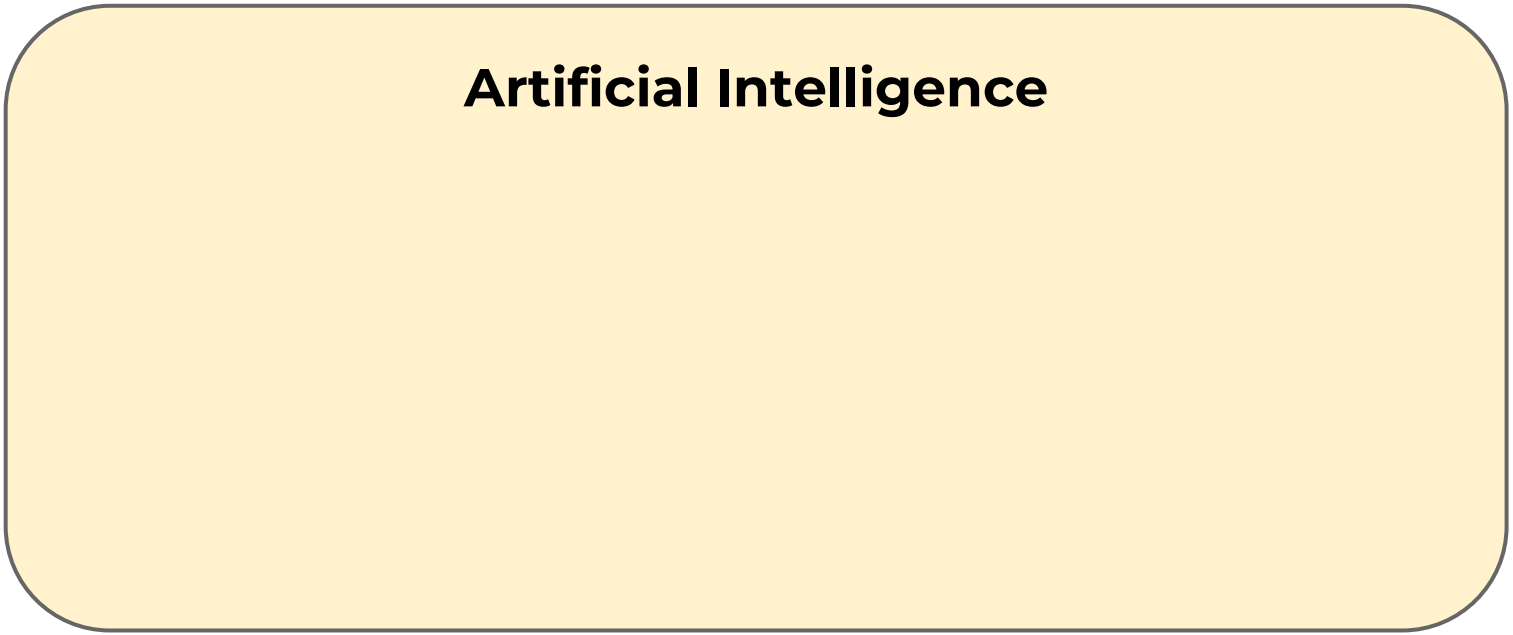

Machine Learning

- Let’s begin by exploring the domains:

Artificial Intelligence

- Intelligence demonstrated by machines.

Machine Learning

- Let’s begin by exploring the domains:

Artificial Intelligence

- Intelligence demonstrated by machines.
  - What is “intelligence”?
  - How to test for “intelligence”?

Machine Learning

- Tests for Artificial Intelligence:
  - Turing Test:

Machine Learning

- Tests for Artificial Intelligence:
  - Turing Test:

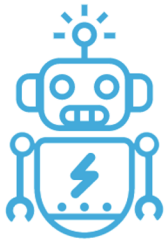

Machine Learning

- Tests for Artificial Intelligence:
  - Turing Test:

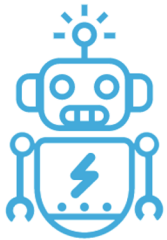

Machine Learning

- Tests for Artificial Intelligence:
  - Turing Test:

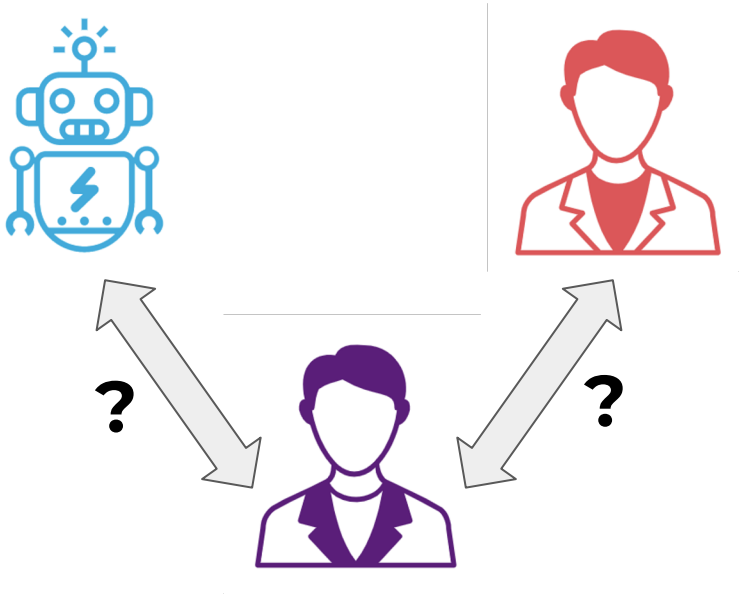

Machine Learning

- Tests for Artificial Intelligence:
  - Marcus Test:
    - Measures a computer’s ability to understand a television program.
  - Lovelace 2.0 Test:
    - Measuring a computer’s ability to create artistic artifacts.

Machine Learning

Artificial Intelligence

- Intelligence demonstrated by machines.
  - What is “intelligence”?
  - How to test for “intelligence”?

Machine Learning

Artificial Intelligence

- General Artificial Intelligence
  - Human level (or better) intelligence in multiple domains.
- Narrow Artificial Intelligence
  - Human level intelligence in a specific domain (e.g. Chatbot, Image Recognition, etc...)

Machine Learning

Artificial Intelligence

- What subdomains are necessary to create artificial intelligence?

Machine Learning

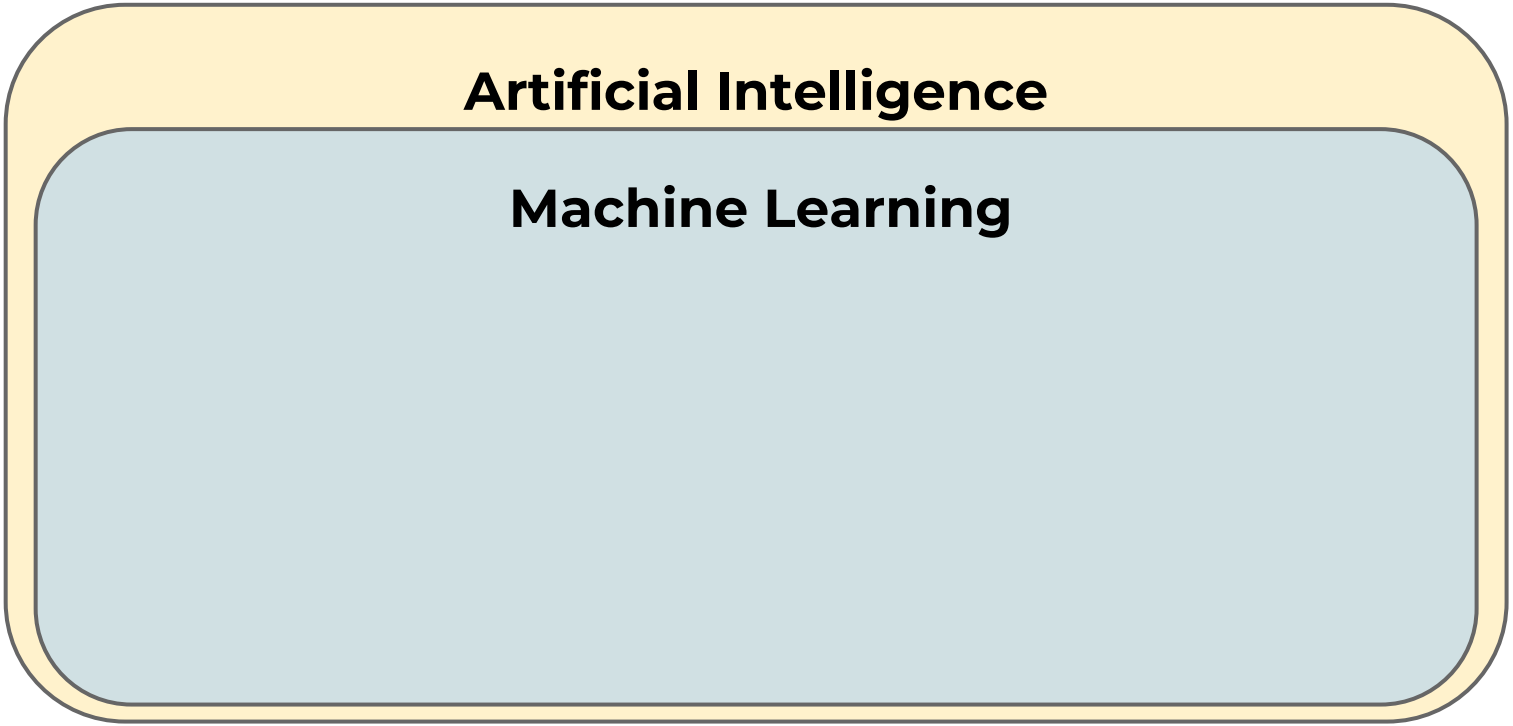

Machine Learning

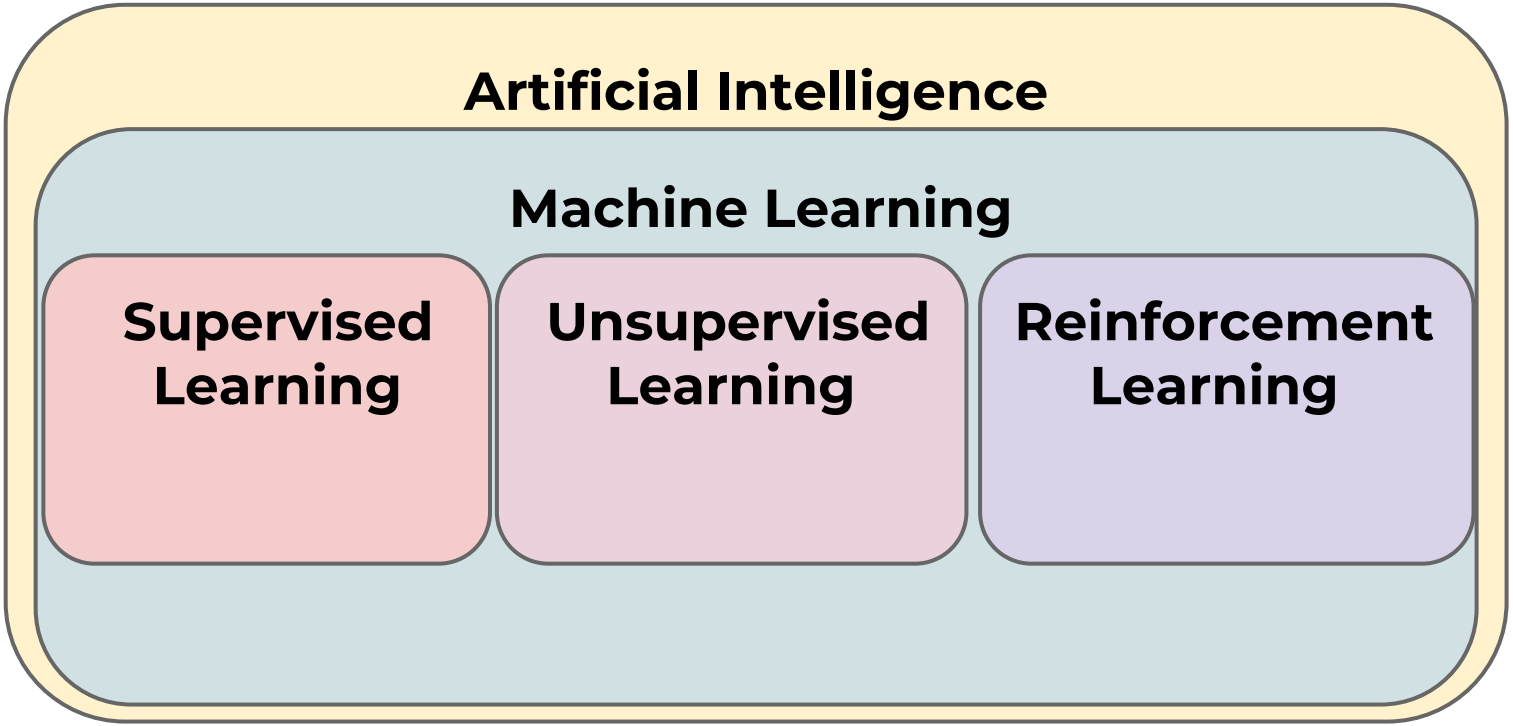

Machine Learning

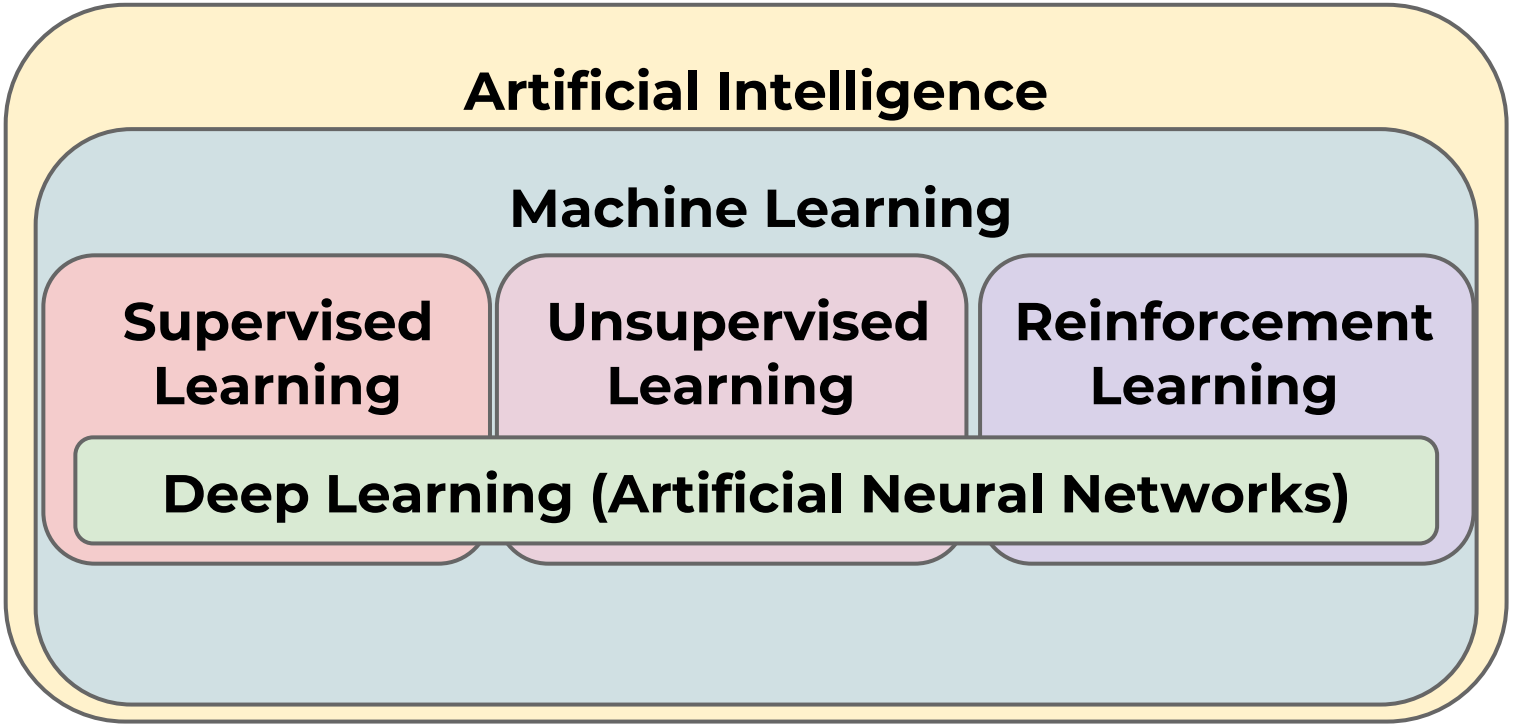

Machine Learning

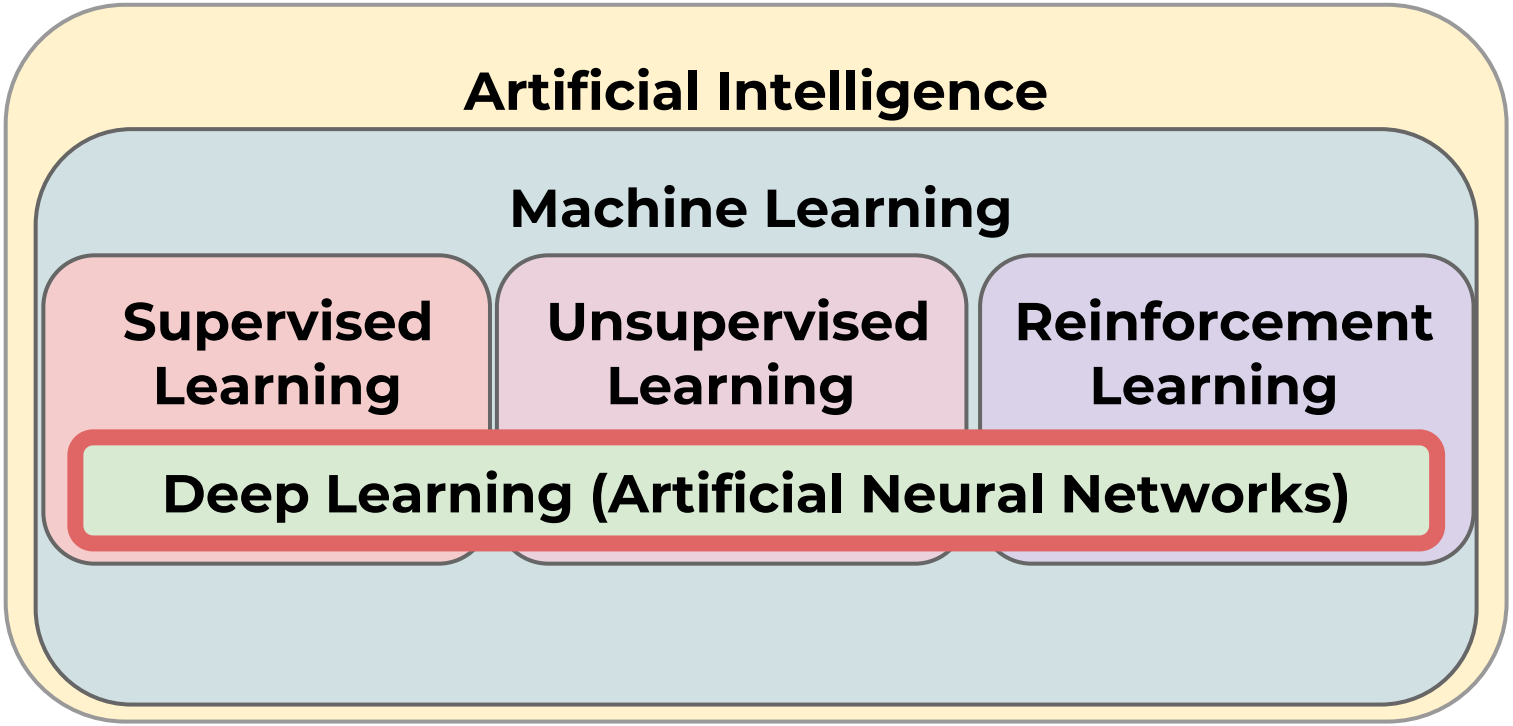

Machine Learning

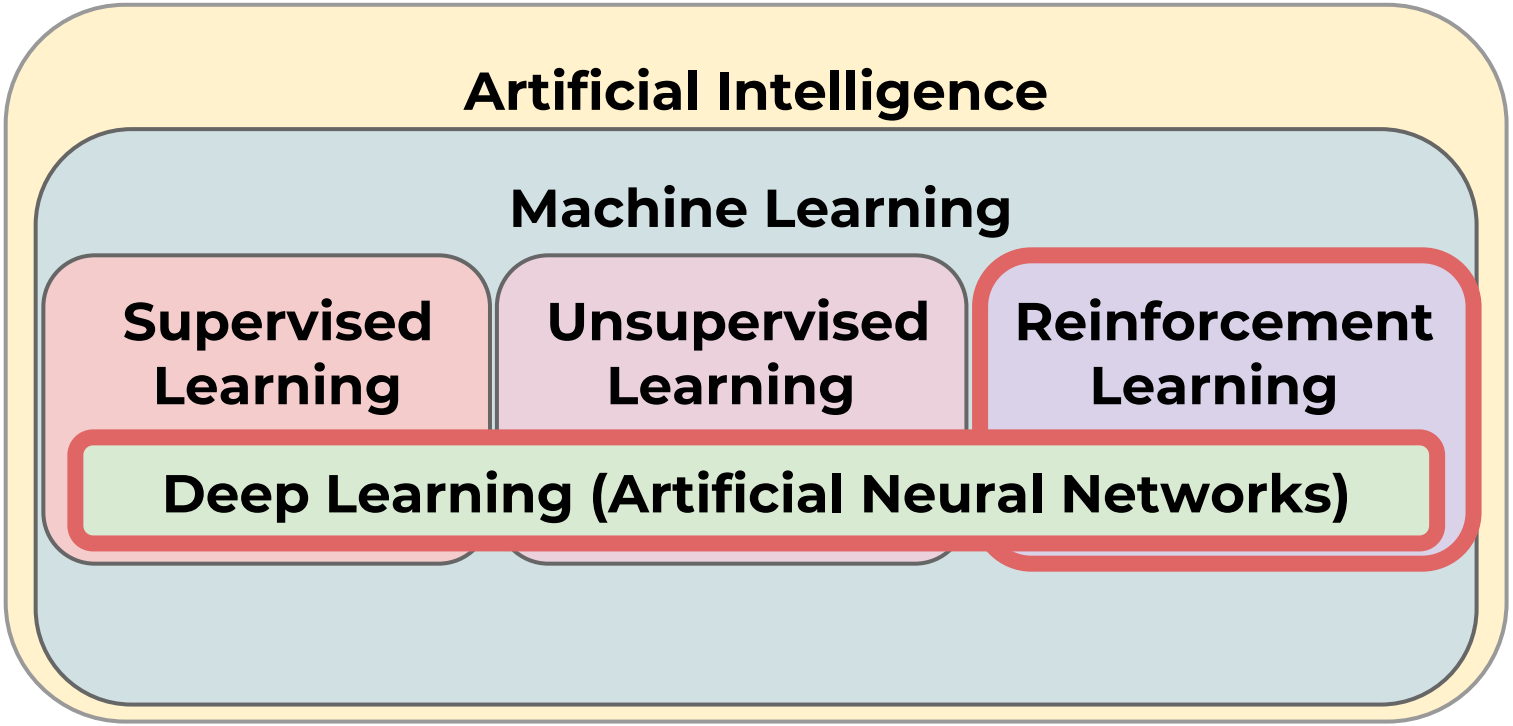

Machine Learning

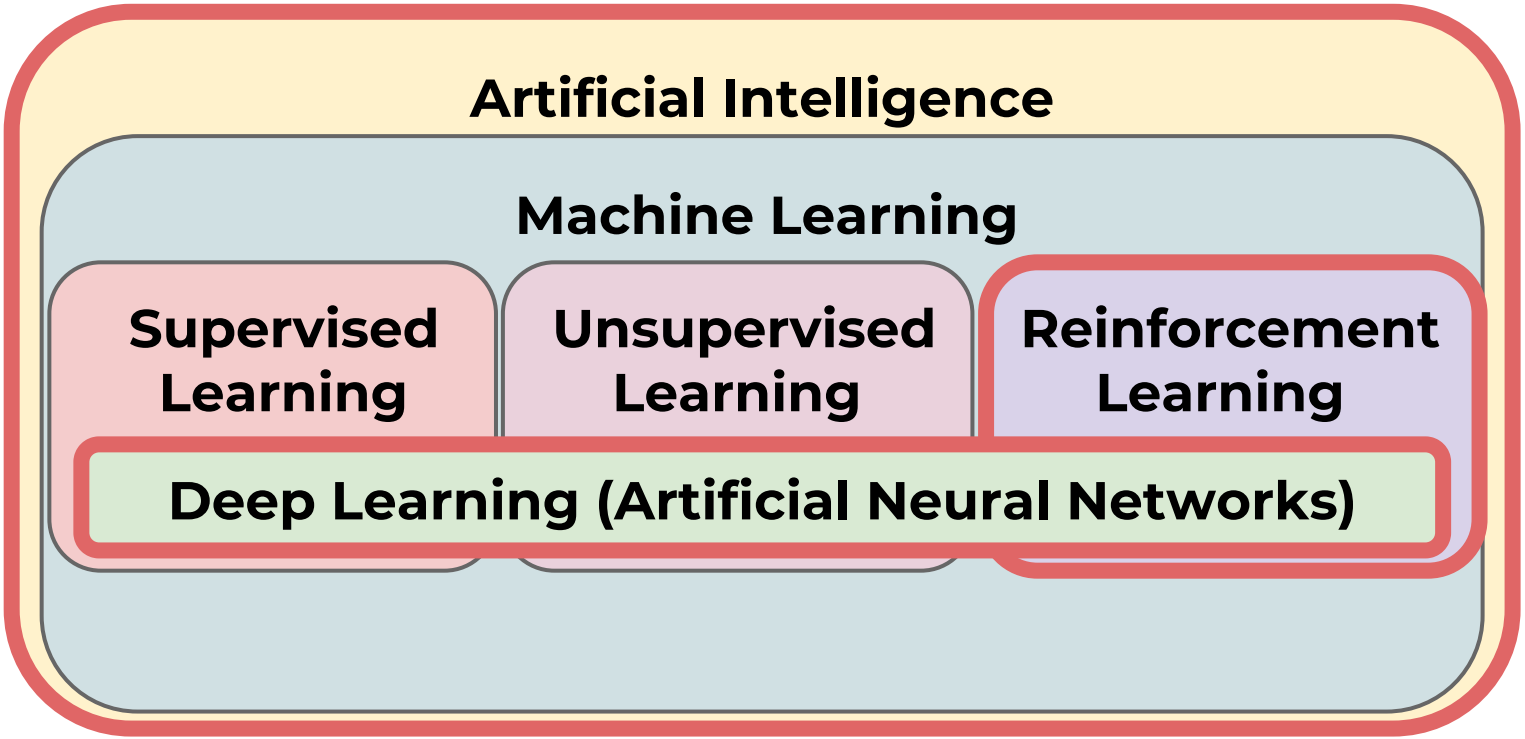

Machine Learning

- Knowledge Path for Artificial Intelligence:

Machine Learning

- Knowledge Path for Artificial Intelligence:
  - Understand Machine Learning
    - Key Library Ideas (Pandas and Scikit-Learn)
    - Supervised Learning Process
    - Deep Learning (ANN and CNN)

Machine Learning

- Knowledge Path for Artificial Intelligence:
  - Understand Reinforcement Learning
    - Agent, Environment, and Policy
    - Tabular Q-Learning

Machine Learning

- Knowledge Path for Artificial Intelligence:
  - Combine Deep Learning and Reinforcement Learning
    - Combine ANN with Q-Learning

Machine Learning

- Knowledge Path for Artificial Intelligence:

Key ML Theory

Machine Learning Concepts

Machine Learning

- Knowledge Path for Artificial Intelligence:

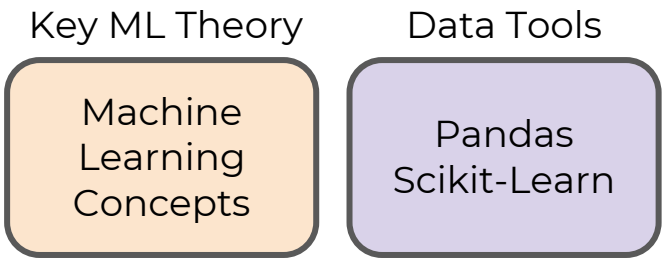

Machine Learning

- Knowledge Path for Artificial Intelligence:

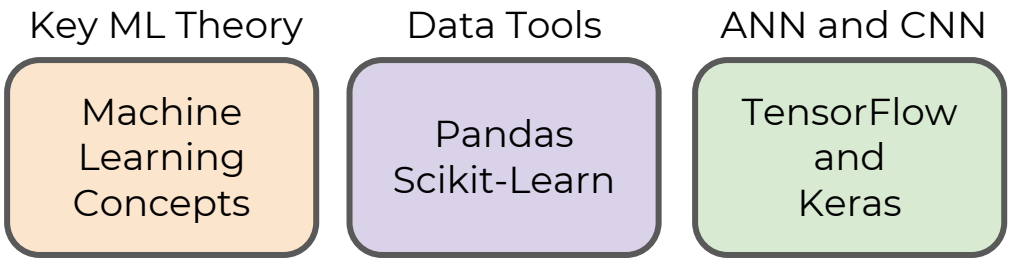

Machine Learning

- Knowledge Path for Artificial Intelligence:

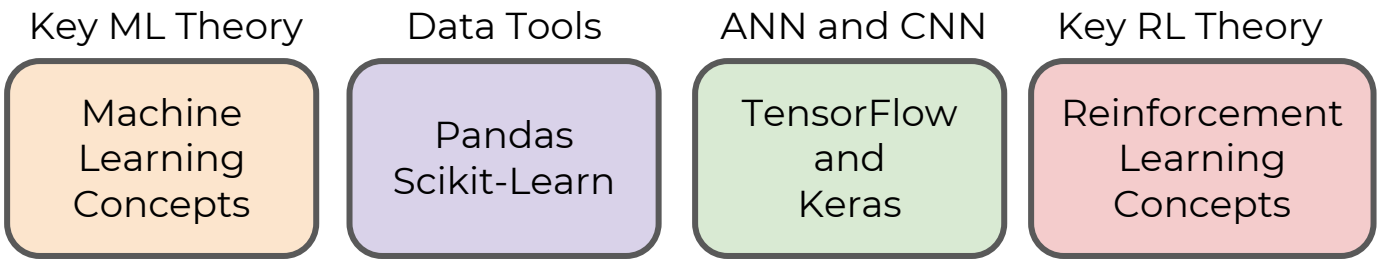

Machine Learning

- Knowledge Path for Artificial Intelligence:

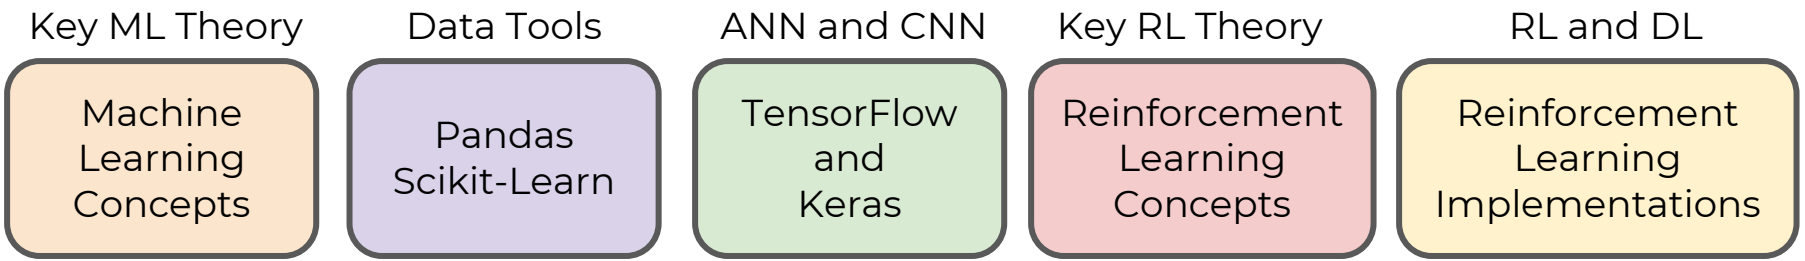

Let’s get started!

Environment Setup

Environment Set-up

- We use a wide variety of libraries in this course, so we'll show you how to set-up a separate virtual environment with Anaconda in order to **pip install** the libraries later on, including **gym** and **tensorflow.**

Environment Set-up

- The easiest way to do this is through the command line.
- Let’s use the:
  - Windows
    - Anaconda Prompt
  - MacOS/Linux
    - Terminal

Environment Set-up

Machine Learning Supervised Pathway

Machine Learning

- Machine learning in general is the study of statistical computer algorithms that improve automatically through data.
- This means unlike typical computer algorithms that rely on human input for what approach to take, ML algorithms infer best approach from the data itself.

Machine Learning

- Machine learning is actually a subset of Artificial Intelligence.
- ML algorithms are not explicitly programmed on which decisions to make.
- Instead the algorithm is designed to infer from the data the most optimal choices to make.

Machine Learning

- What kinds of problems can ML solve?
  - Credit Scoring
  - Insurance Risk
  - Price Forecasting
  - Spam Filtering
  - Customer Segmentation
  - Much more!

“Standard” ML Pathway

Problem to Solve

Real World

Question to Answer

ML Pathway

Real World

Problem to Solve

How to fix or change X?

Question to Answer

How does a change in X affect Y?

ML Pathway

Real World

Problem to Solve

How to fix or change X?

Question to Answer

How does a change in X affect Y?

ML Pathway

Real World

Data Product

Data Analysis

ML Pathway

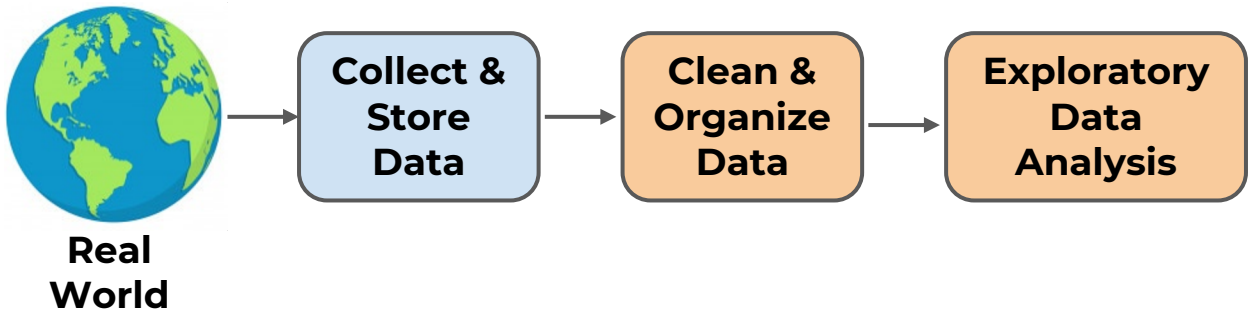
Real World
Collect & Store Data
Clean & Organize Data
Exploratory Data Analysis

ML Pathway

Real World

Collect & Store Data → **Clean & Organize Data** → **Exploratory Data Analysis** → **Machine Learning Models**

Supervised Learning:
Predict an Outcome
Unsupervised Learning:
Discover Patterns in Data

ML Pathway

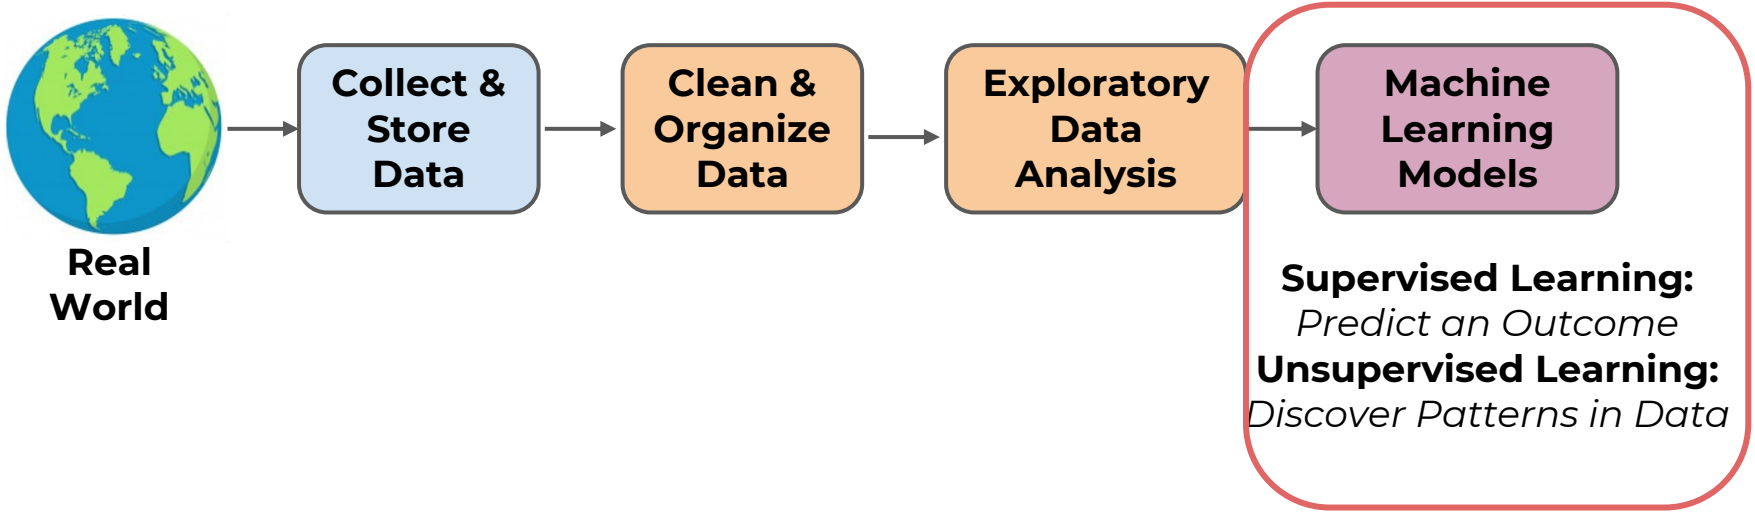
Real World
Collect & Store Data
Clean & Organize Data
Exploratory Data Analysis
Machine Learning Models
Supervised Learning:
Predict an Outcome
Unsupervised Learning:
Discover Patterns in Data

Why Machine Learning?

Machine Learning

- Structure of ML Problem framing:
  - Given **features** from a data set **obtain** a desired **label**.
  - ML algorithms are often called “estimators” since they are estimating the desired **label** or output.

Machine Learning

- How can ML be so robust in solving all sorts of problems?
- Machine learning algorithms rely on data and a set of statistical methods to learn what features are important in data.

Machine Learning

- Simple Example:
  - Predict the price a house should sell at given its current features (Area,Bedrooms,Bathrooms,etc...)

Machine Learning

- House Price Prediction
  - Typical Algorithm
    - Human user defines an algorithm to manually set values of importance for each feature.

Machine Learning

- House Price Prediction
  - ML Algorithm
    - Algorithm automatically determines importance of each feature from existing data

Machine Learning

- Why machine learning?
  - Many complex problems are only solvable with machine learning techniques.
  - Problems such as spam email or handwriting identification require ML for an effective solution.

Machine Learning

- Why not just use machine learning for everything?
  - Major caveat to effective ML is good data.
  - Majority of development time is spent cleaning and organizing data, **not** implementing ML algorithms.

Machine Learning

- Do we develop our own ML algorithms?
  - Rare to have a need to manually develop and implement a new ML algorithm, since these techniques are well documented and developed.

Machine Learning

- Let’s continue this discussion by exploring the types of machine learning algorithms!

Types of Machine Learning

Machine Learning

- There are three main types of Machine Learning:
  - Supervised Learning
  - Unsupervised Learning
  - Reinforcement Learning

Machine Learning

- Supervised Learning
  - Using **historical** and **labeled** data, the machine learning model predicts a value.
- Unsupervised Learning
  - Applied to **unlabeled** data, the machine learning model discovers possible patterns in the data.

Machine Learning

- Supervised Learning
  - Requires **historical labeled** data:
    - Historical
      - Known results and data from the past.
    - Labeled
      - The desired output is known.

Machine Learning

- Supervised Learning
  - Two main label types
    - Categorical Value to Predict
      - Classification Task
    - Continuous Value to Predict
      - Regression Task

Machine Learning

- Supervised Learning
  - Classification Tasks
    - Predict an assigned category
      - Cancerous vs. Benign Tumor
      - Fulfillment vs. Credit Default
      - Assigning Image Category
        - Handwriting Recognition

Machine Learning

- Supervised Learning
  - Regression Tasks
    - Predict a continuous value
      - Future prices
      - Electricity loads
      - Test scores

Machine Learning

- Unsupervised Learning
  - Group and interpret data without a historical label.
  - Example:
    - Clustering customers into separate groups based off their behaviour features.

Machine Learning

- Unsupervised Learning
  - Major downside is because there was no historical “correct” label, it is much harder to evaluate performance of an unsupervised learning algorithm.

Machine Learning

- Reinforcement Learning

Machine Learning

- Finally, before we dive into coding and linear regression in the next section, let’s have a deep dive into the entire Supervised Machine Learning process to set ourselves up for success!

Supervised Machine Learning Process

Machine Learning

- Machine Learning Pathway

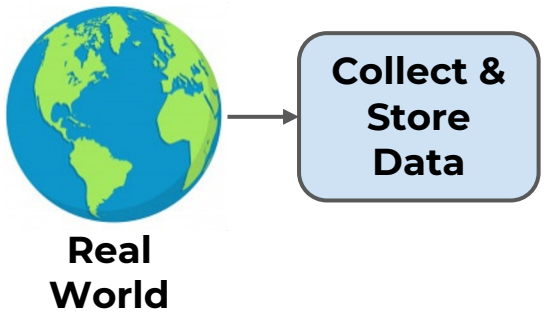

Machine Learning

- Machine Learning Pathway

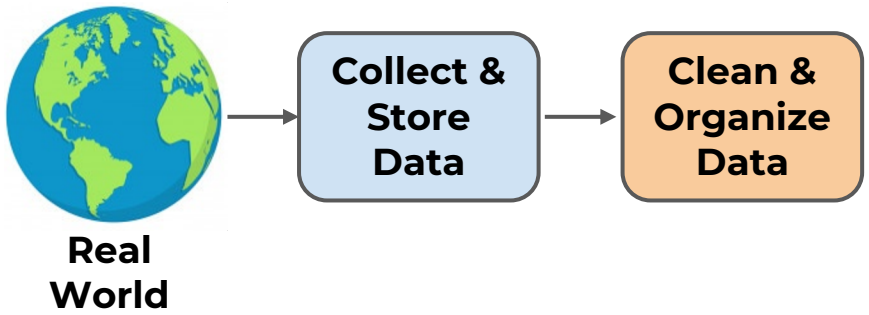

Machine Learning

- Machine Learning Pathway

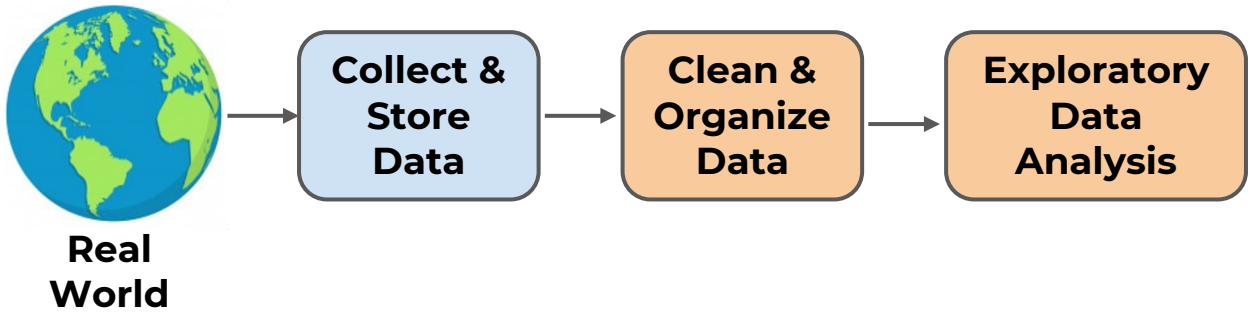

Machine Learning

- Machine Learning Pathway

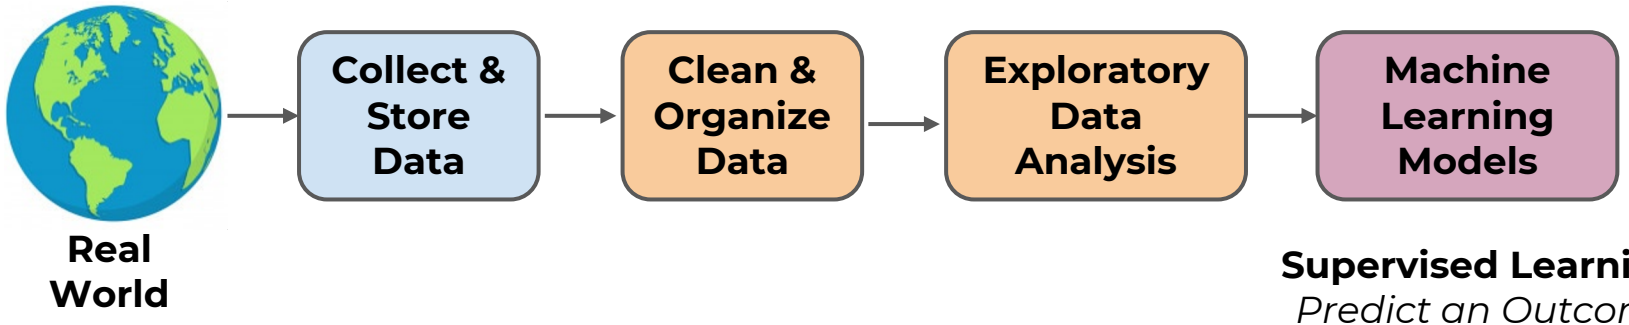

Supervised Learning:
Predict an Outcome
Unsupervised Learning:
Discover Patterns in Data

Machine Learning

- Machine Learning Pathway

Real World

Collect & Store Data → **Clean & Organize Data** → **Exploratory Data Analysis** → **Machine Learning Models**

Jupyter,NumPy, Pandas, Matplotlib, Seaborn

Supervised Learning:
Predict an Outcome
Unsupervised Learning:
Discover Patterns in Data

Scikit-learn or TensorFlow

Machine Learning

- Machine Learning Pathway

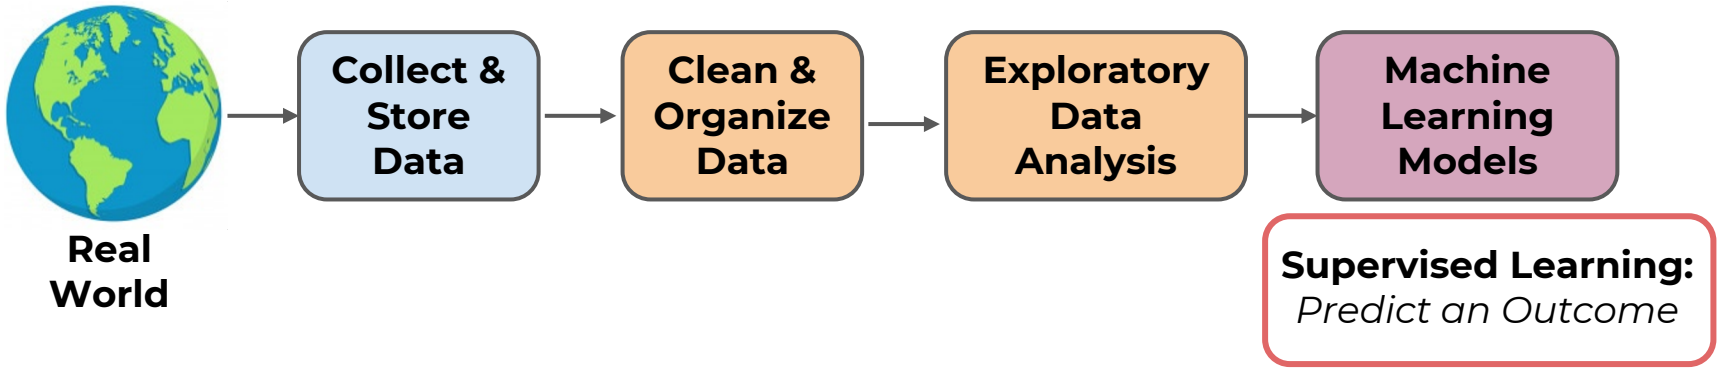

Machine Learning

- ML Process : Supervised Learning Tasks

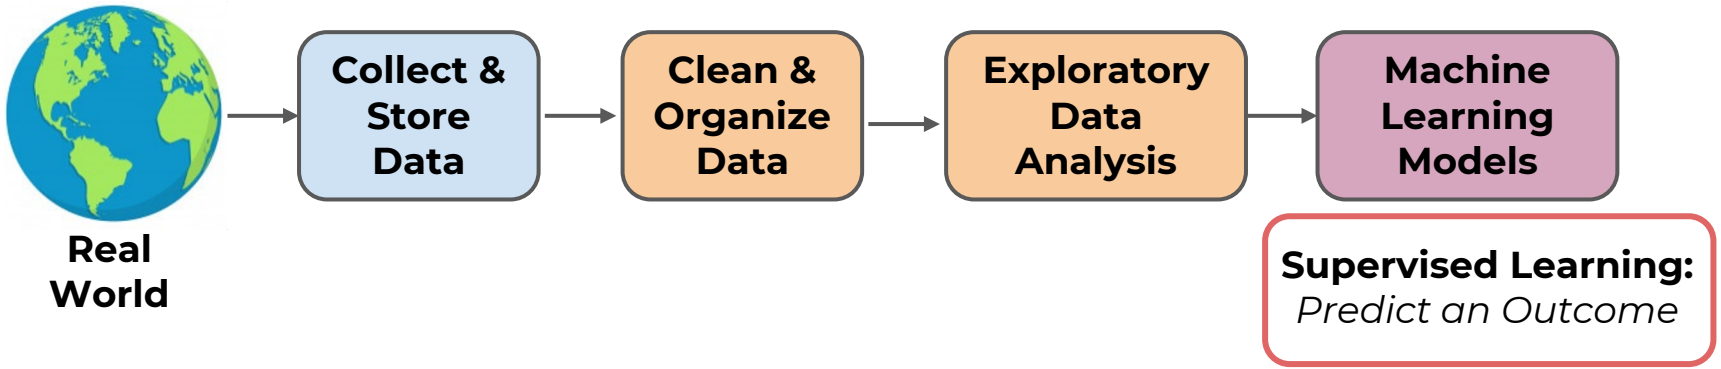

Machine Learning

- **Predict price a house should sell at.**

Real World

Collect & Store Data → **Clean & Organize Data** → **Exploratory Data Analysis** → **Machine Learning Models**

Supervised Learning:
Predict an Outcome

Machine Learning

- **Supervised** Machine Learning Process
- Start with collecting and organizing a data set based on history:

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Machine Learning

- **Historical labeled** data on previously sold houses.

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Machine Learning

- If a new house comes on the market with a known Area, Bedrooms, and Bathrooms: *Predict what price should it sell at.*

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Machine Learning

- Data Product:
  - Input house features
  - Output predicted selling price

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Machine Learning

- Using **historical, labeled** data predict a future outcome or result.

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Machine Learning

- **Predict price a house should sell at.**

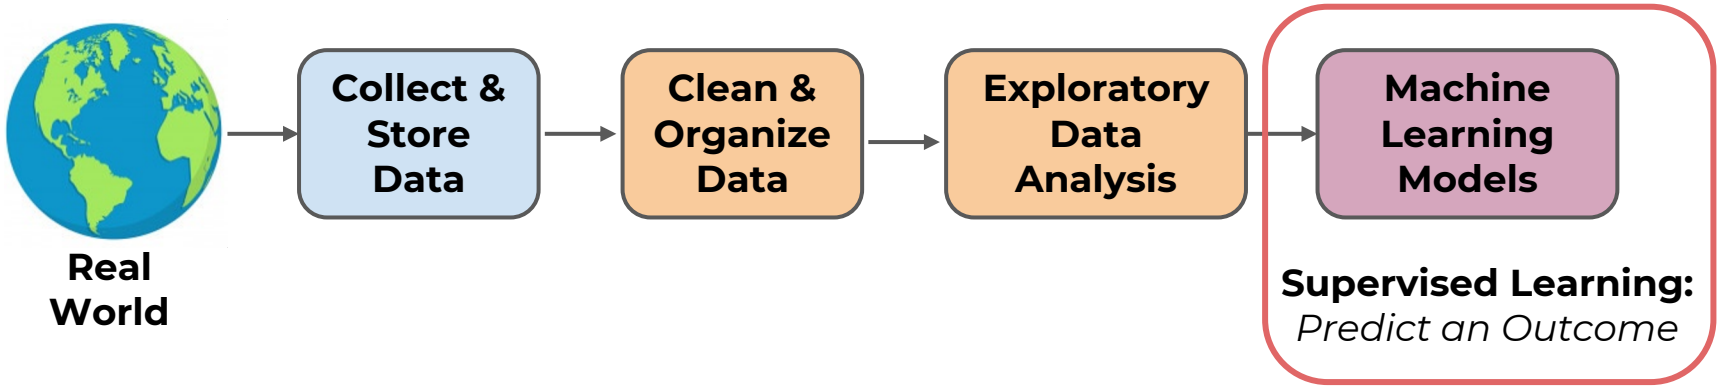

Machine Learning

- **Predict price a house should sell at.**

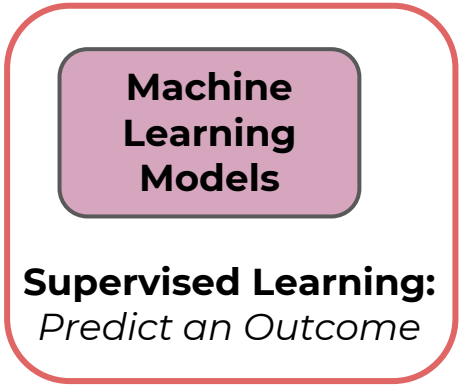

Machine Learning

- **Predict price a house should sell at.**

Machine Learning Models

Supervised Learning:
Predict an Outcome

Machine Learning

- **Predict price a house should sell at.**

Machine Learning Models

Supervised Learning:
Predict an Outcome

Data

Machine Learning

- **Supervised** Machine Learning Process

Machine Learning

- **Supervised** Machine Learning Process

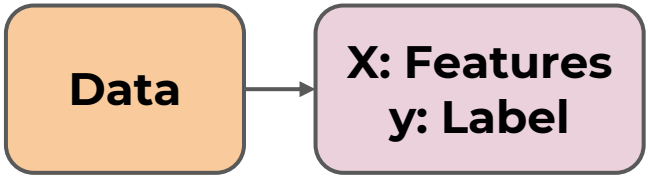

Machine Learning

- **Supervised** Machine Learning Process

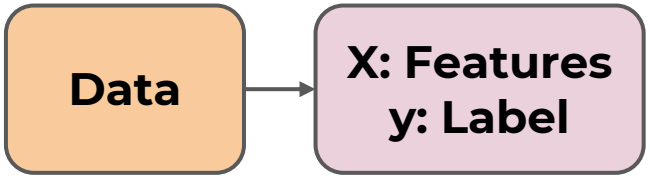

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Machine Learning

- **Label** is what we are trying to predict

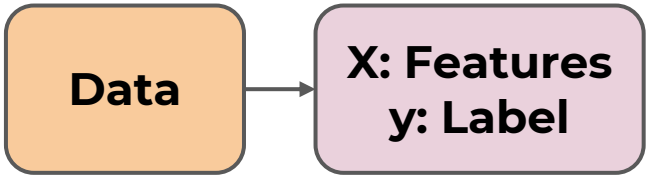

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Machine Learning

- **Label** is what we are trying to predict

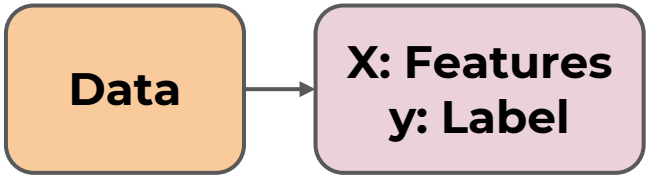

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Machine Learning

- **Features** are known characteristics or components in the data

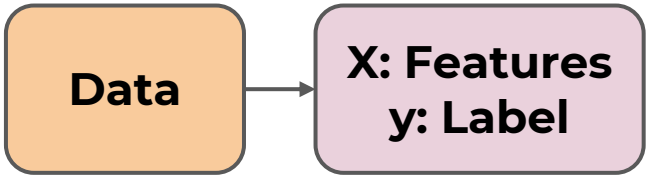

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Machine Learning

- **Features** are known characteristics or components in the data

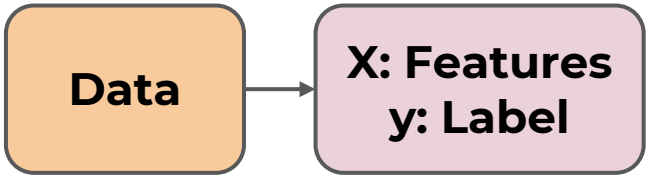

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Machine Learning

- **Features** and **Label** are identified according to the problem being solved.

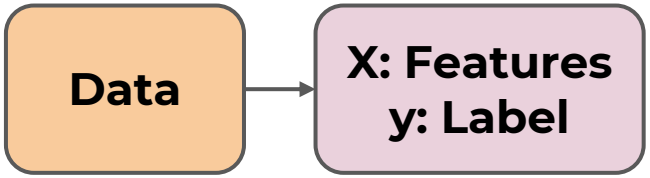

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- Split data into training set and test set

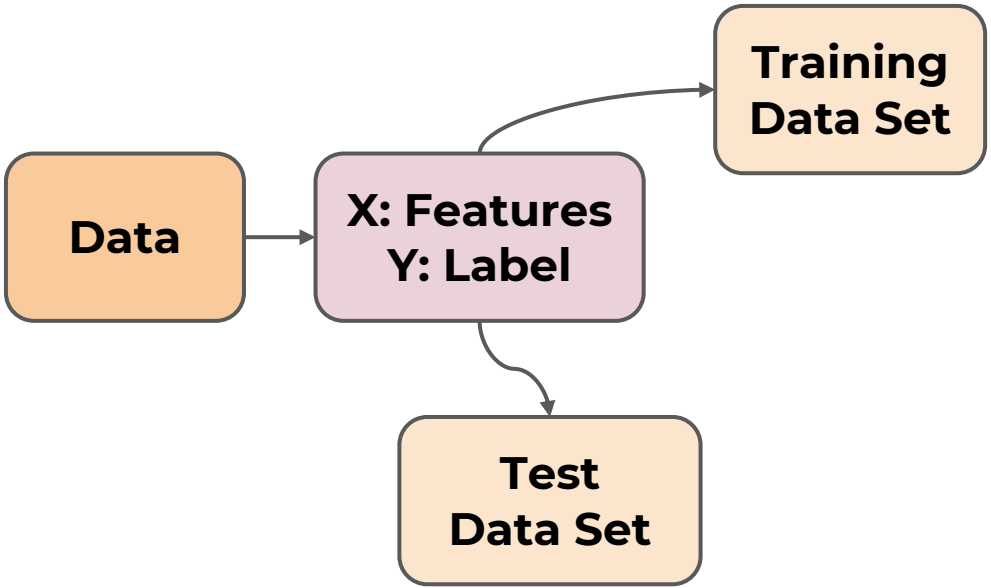

Supervised Machine Learning Process

- Later on we will discuss cross-validation

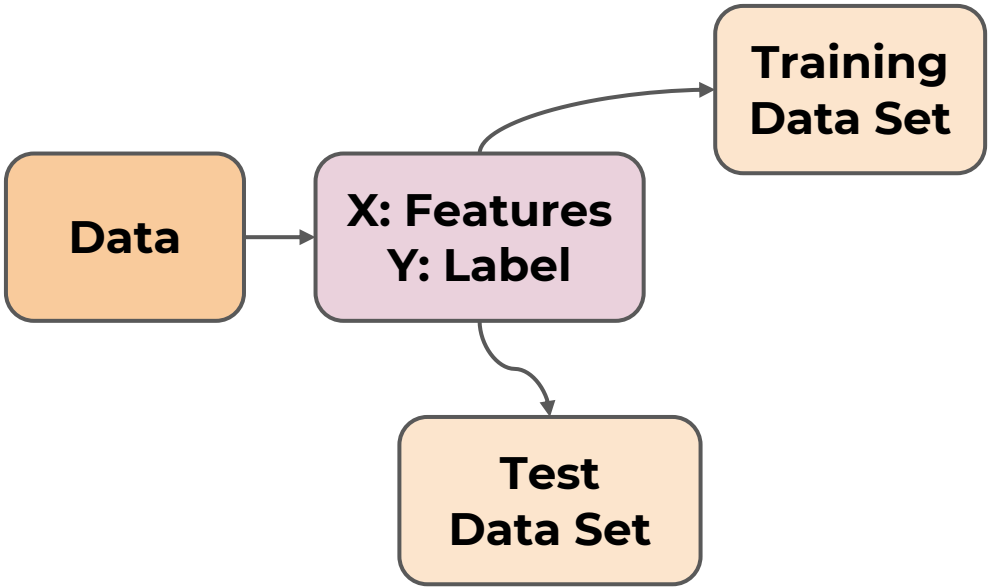

Supervised Machine Learning Process

- Why perform this split? How to split?

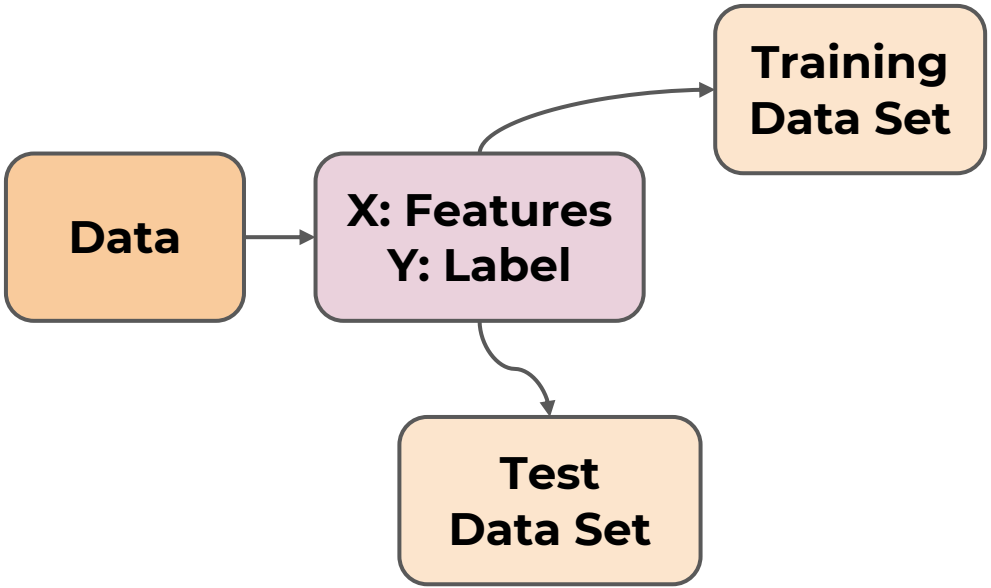

Supervised Machine Learning Process

- Why perform this split? How to split?

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- How would you judge a human realtor's performance?

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- Ask a human realtor to take a look at historical data...

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- Then give her the features of a house and ask her to predict a selling price.

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- But how would you measure how accurate her prediction is? What house should you choose to test on?

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- You can’t judge her based on a new house that hasn’t sold yet, you don’t know it’s true selling price!

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- You shouldn’t judge her on data she’s already seen, she could have **memorized** it!

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- Thus the need for a Train/Test split of the data, let's explore further...

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- We already organized the data into Features (X) and a Label (y)

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- Now we will split this into a training set and a test set:

TRAIN

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- Now we will split this into a training set and a test set:

	Area m^2	Bedrooms	Bathrooms	Price
TRAIN	200	3	2	$500,000
	190	2	1	$450,000
	230	3	3	$650,000
TEST	180	1	1	$400,000
	210	2	2	$550,000

Supervised Machine Learning Process

- Notice how we have 4 components

Supervised Machine Learning Process

- Let’s go back to fairly testing our human realtor….

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- Let's go back to fairly testing our human realtor....

TRAIN

TEST

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000
180	1	1	$400,000
210	2	2	$550,000

Supervised Machine Learning Process

- Let her study and learn on the training set getting access to both X and y.

TRAIN

Area m^2	Bedrooms	Bathrooms	Price
200	3	2	$500,000
190	2	1	$450,000
230	3	3	$650,000

Supervised Machine Learning Process

- After she has “learned” about the data, we can test her skill on the test set.

TEST

Area m^2	Bedrooms	Bathrooms
180	1	1
210	2	2

Supervised Machine Learning Process

- Provide only the X test data and ask for her predictions for the sell price.

TEST

Area m^2	Bedrooms	Bathrooms
180	1	1
210	2	2

Supervised Machine Learning Process

- This is new data she has never seen before! She has also never seen the real sold price.

TEST

Area m^2	Bedrooms	Bathrooms
180	1	1
210	2	2

Supervised Machine Learning Process

- Ask for predictions per data point.

Predictions	Area m^2	Bedrooms	Bathrooms
$410,000	180	1	1
$540,000	210	2	2

Supervised Machine Learning Process

- Then bring back the original prices.

Predictions	Area m^2	Bedrooms	Bathrooms	Price
$410,000	180	1	1	$400,000
$540,000	210	2	2	$550,000

Supervised Machine Learning Process

- Finally compare predictions against true test price.

Predictions	Price
$410,000	$400,000
$540,000	$550,000

Supervised Machine Learning Process

- This is often labeled as **ŷ** compared again **y**

ŷ Predictions	**y** Price
$410,000	$400,000
$540,000	$550,000

Supervised Machine Learning Process

- Later on we will discuss the many methods of evaluating this performance!

Predictions	Price
$410,000	$400,000
$540,000	$550,000

Supervised Machine Learning Process

- Split Data

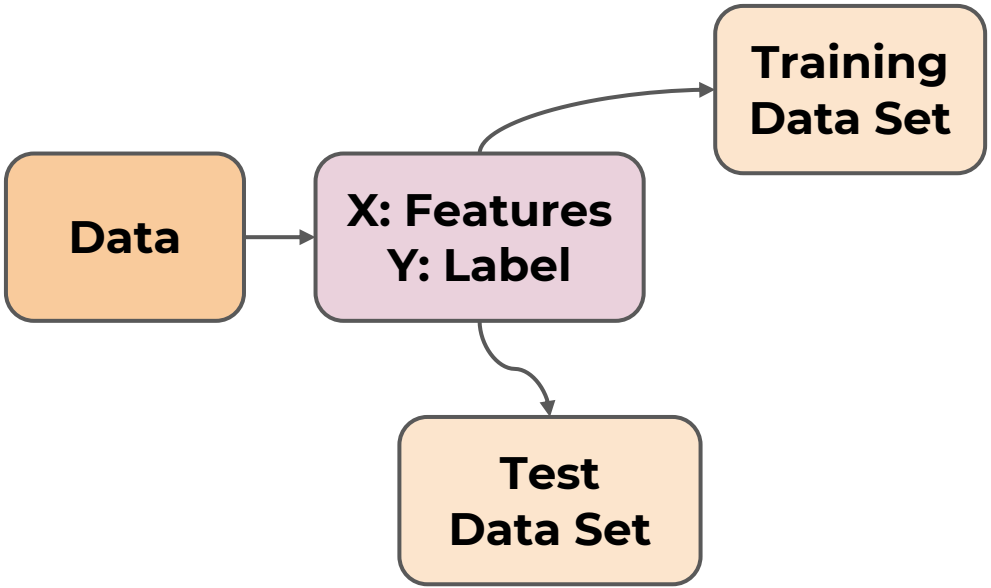

Supervised Machine Learning Process

- Split Data, Fit on Train Data

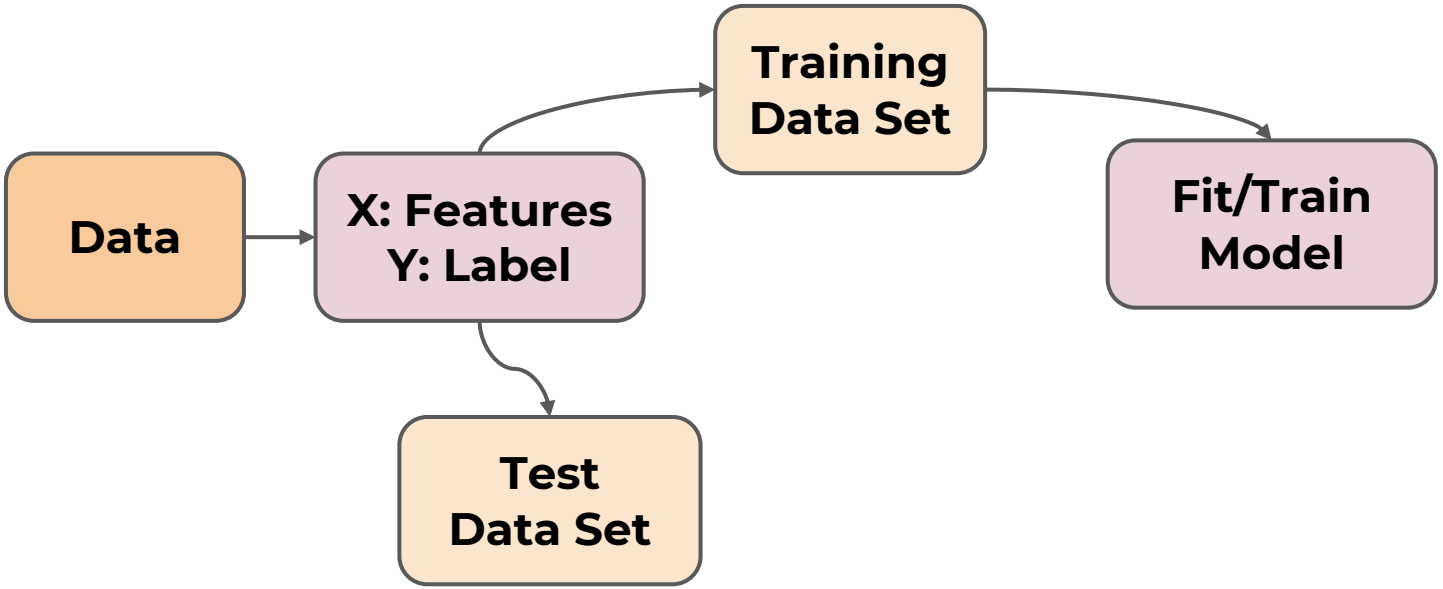

Supervised Machine Learning Process

- Split Data, Fit on Train Data,Evaluate Model

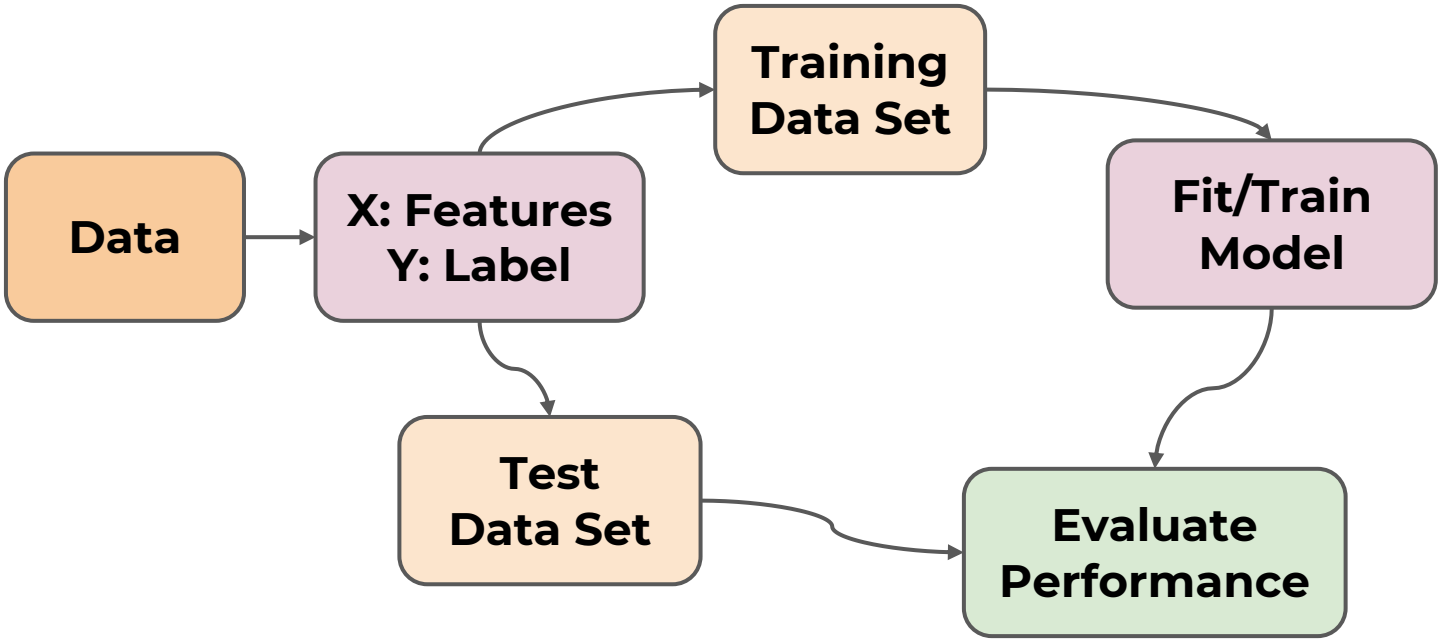

Supervised Machine Learning Process

- What happens if performance isn't great?

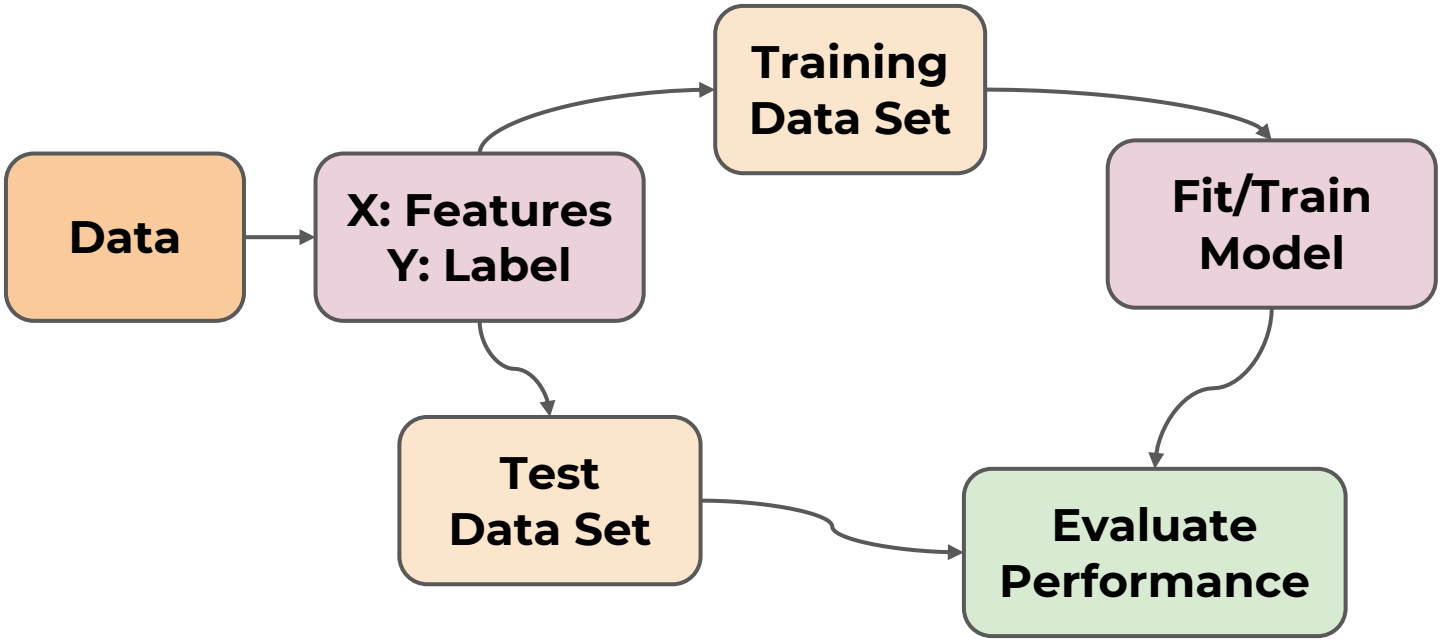

Supervised Machine Learning Process

- We can adjust model **hyperparameters**

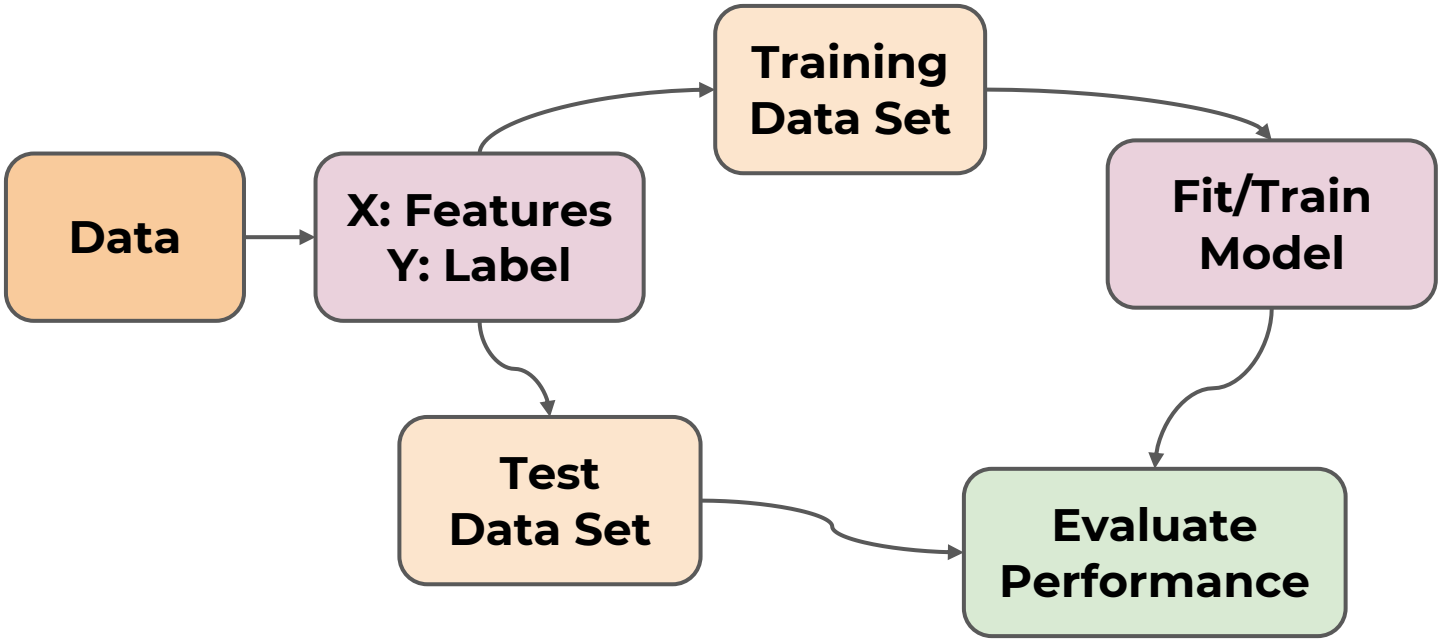

Supervised Machine Learning Process

- Many algorithms have adjustable values

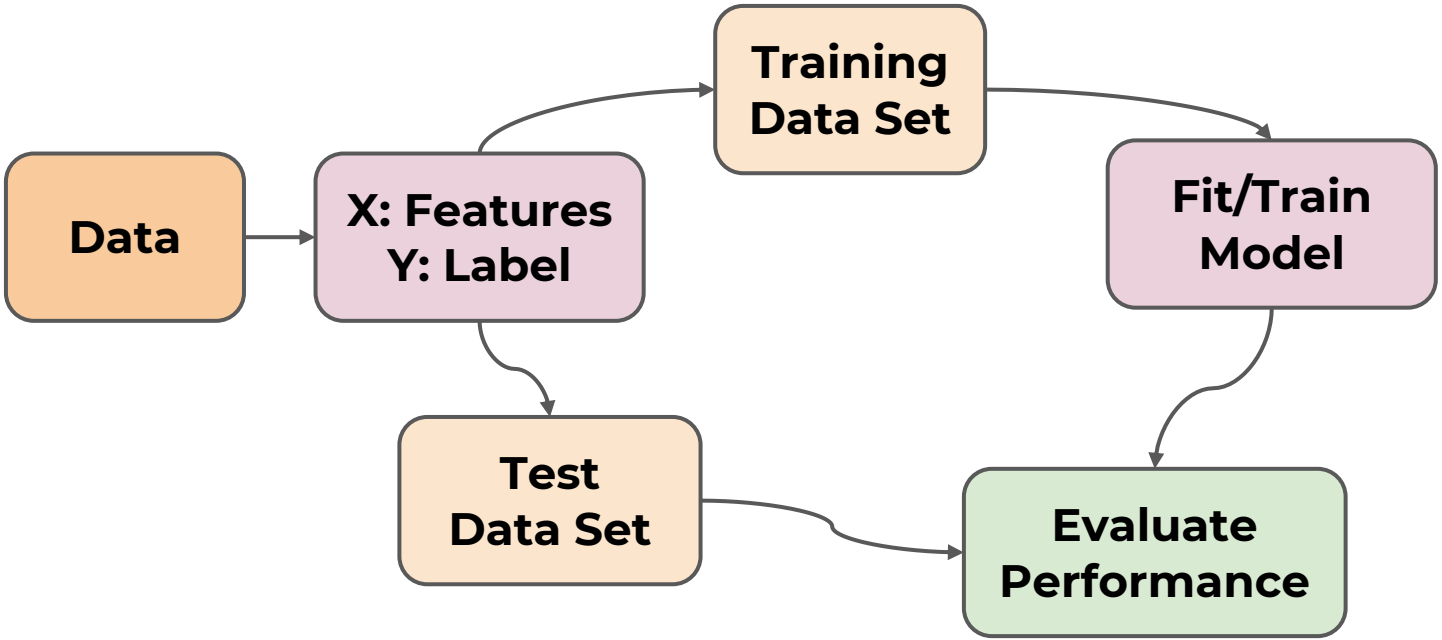

Supervised Machine Learning Process

- Many algorithms have adjustable values

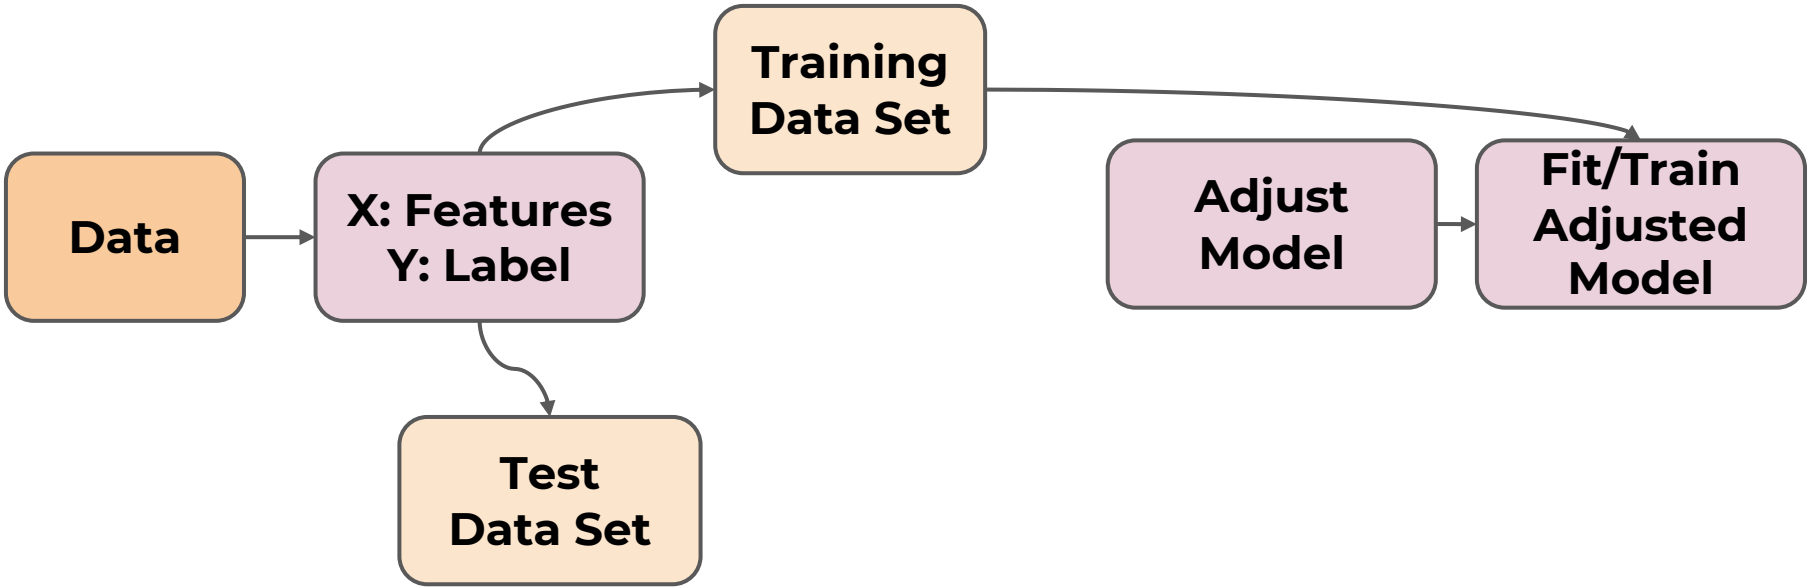

Supervised Machine Learning Process

- Evaluate adjusted model

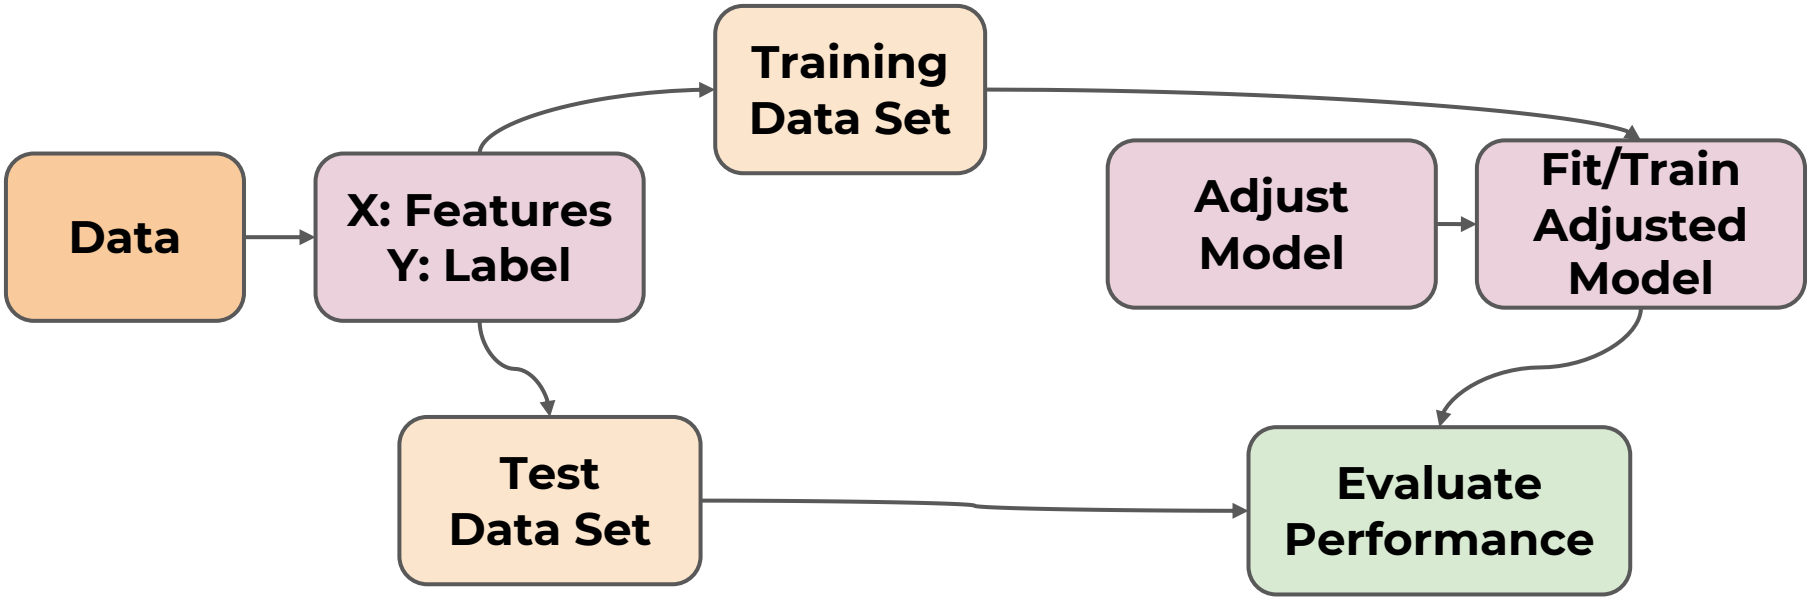

Supervised Machine Learning Process

- Can repeat this process as necessary

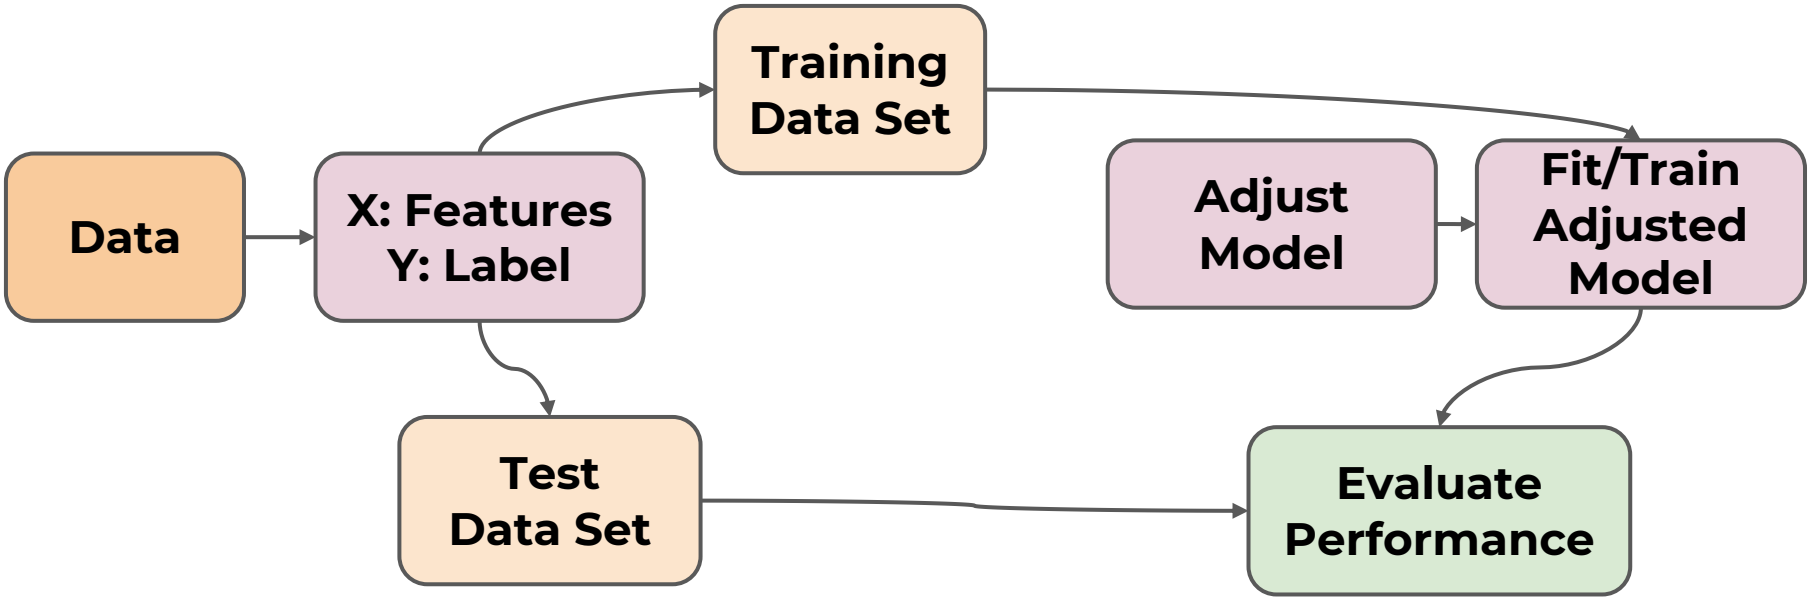

Supervised Machine Learning Process

- Full and Simplified Process

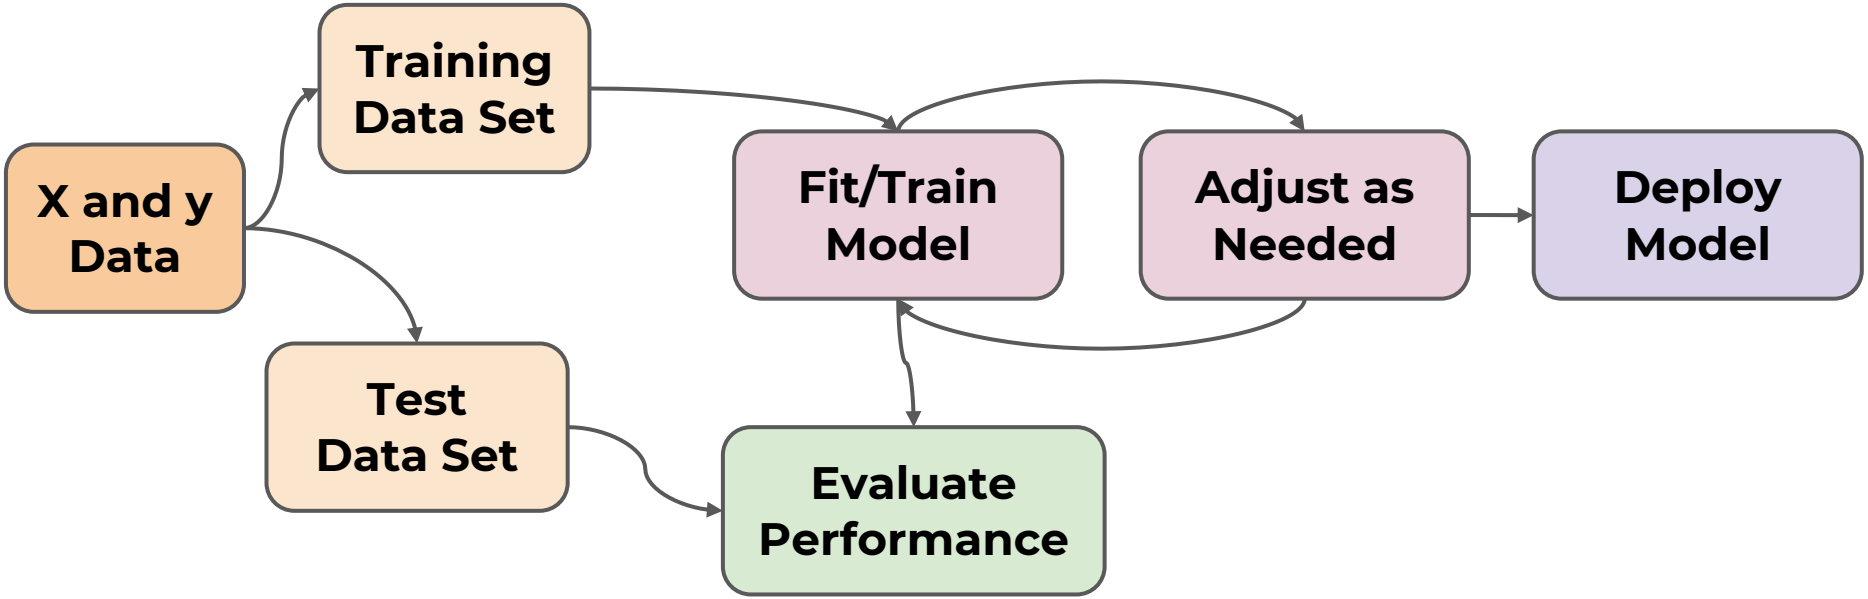

Supervised Machine Learning Process

- Get X and y data

X and y Data

Supervised Machine Learning Process

- Split data for evaluation purposes

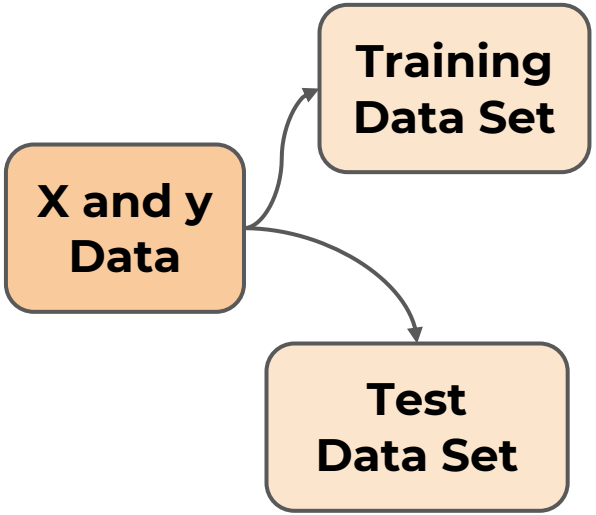

Supervised Machine Learning Process

- Fit ML Model on Training Data Set

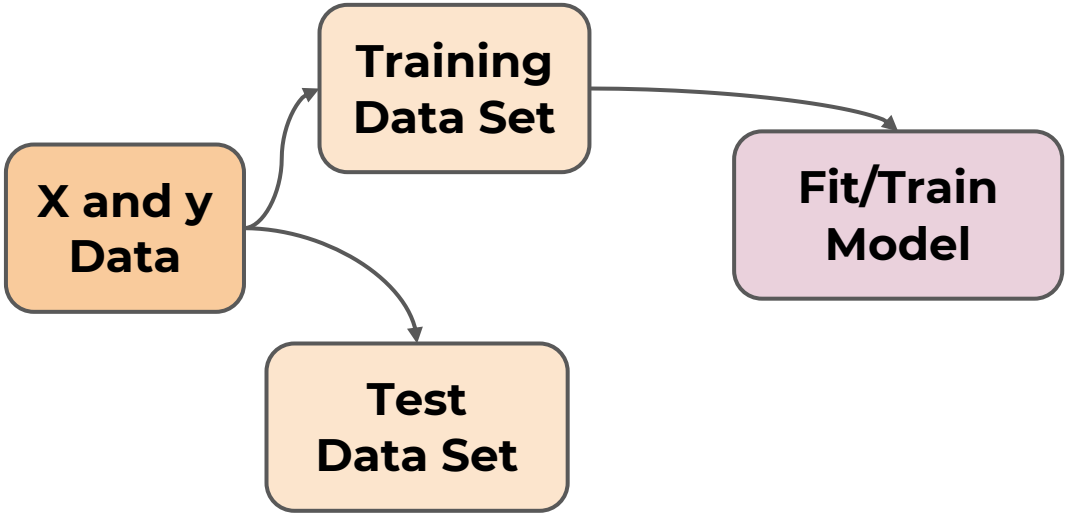

Supervised Machine Learning Process

- Evaluate Model Performance

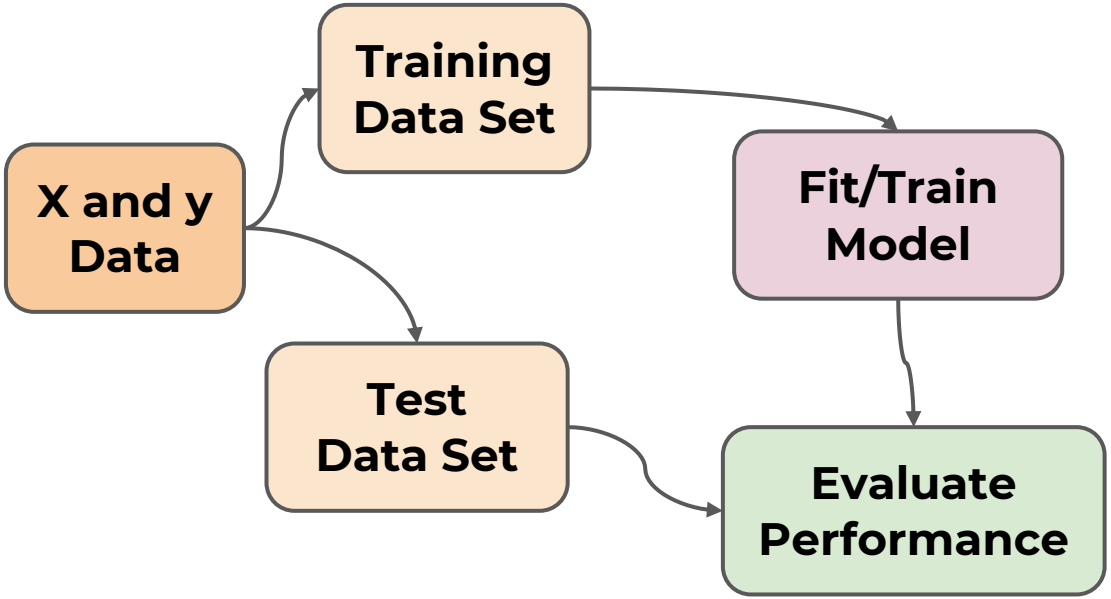

Supervised Machine Learning Process

- Adjust model hyperparameters as needed

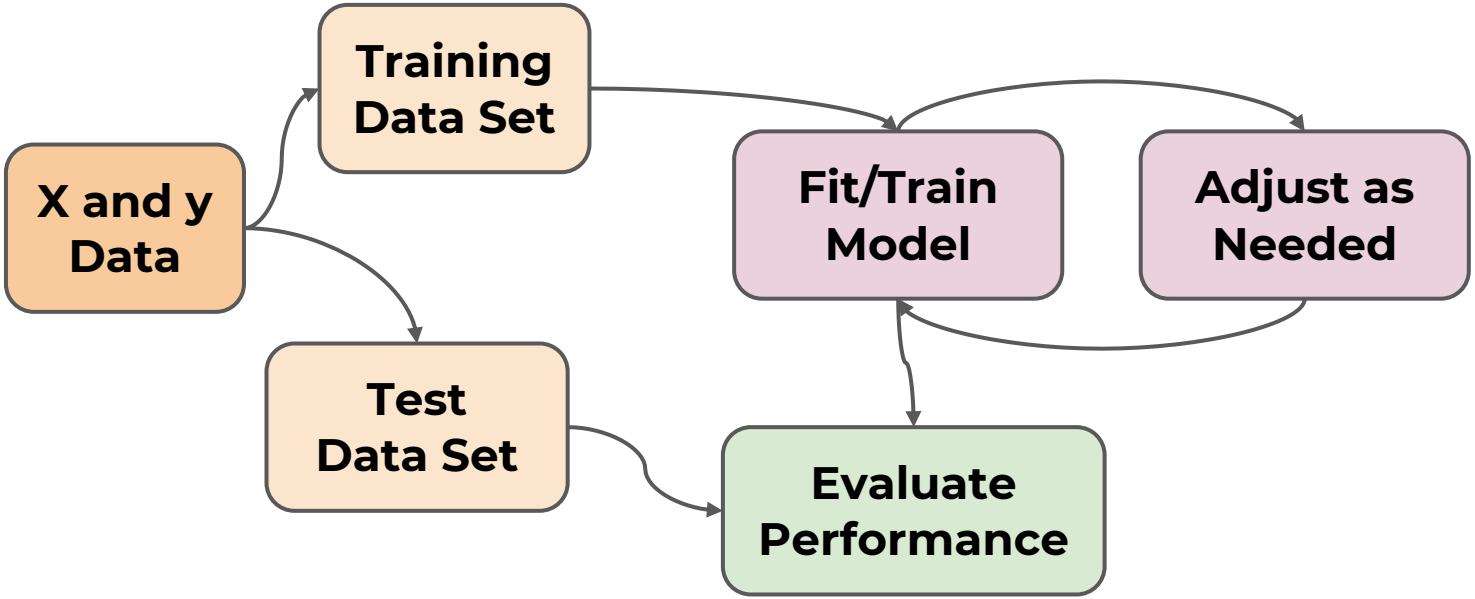

Supervised Machine Learning Process

- Deploy model to real world

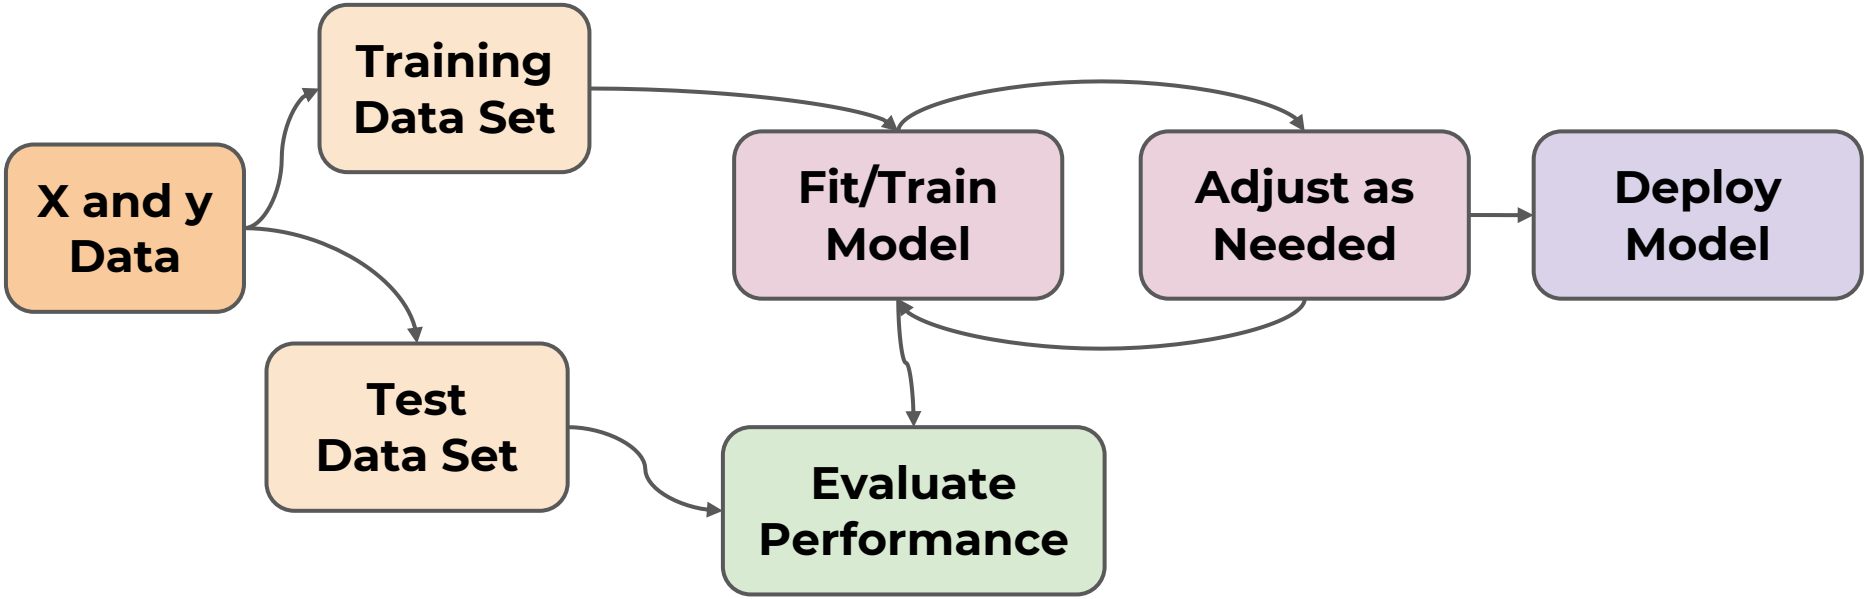

Machine Learning

- ML Process : Supervised Learning Tasks

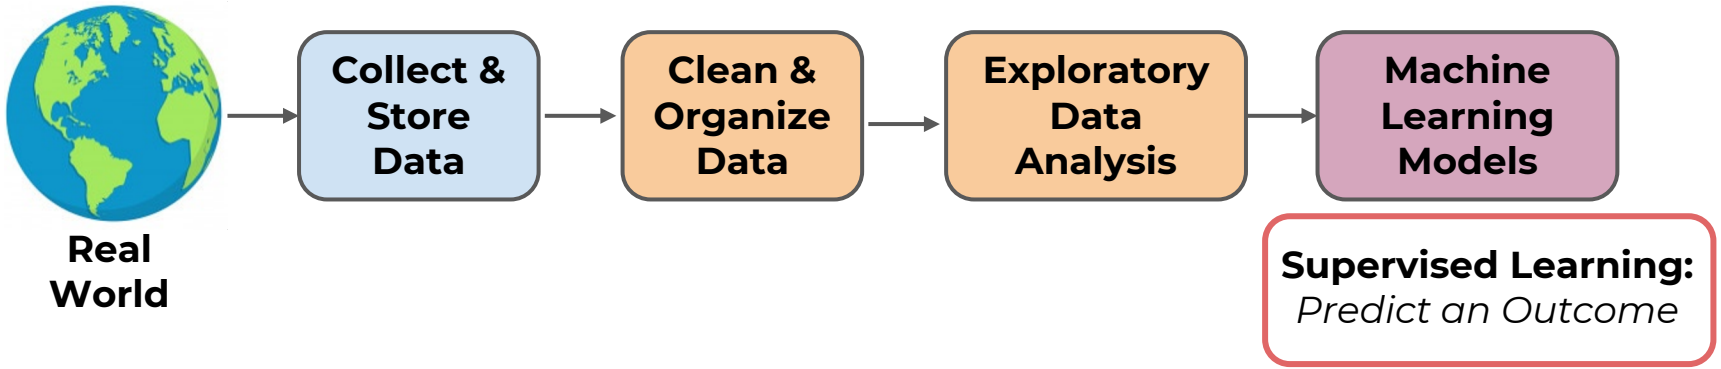

ML Pathway

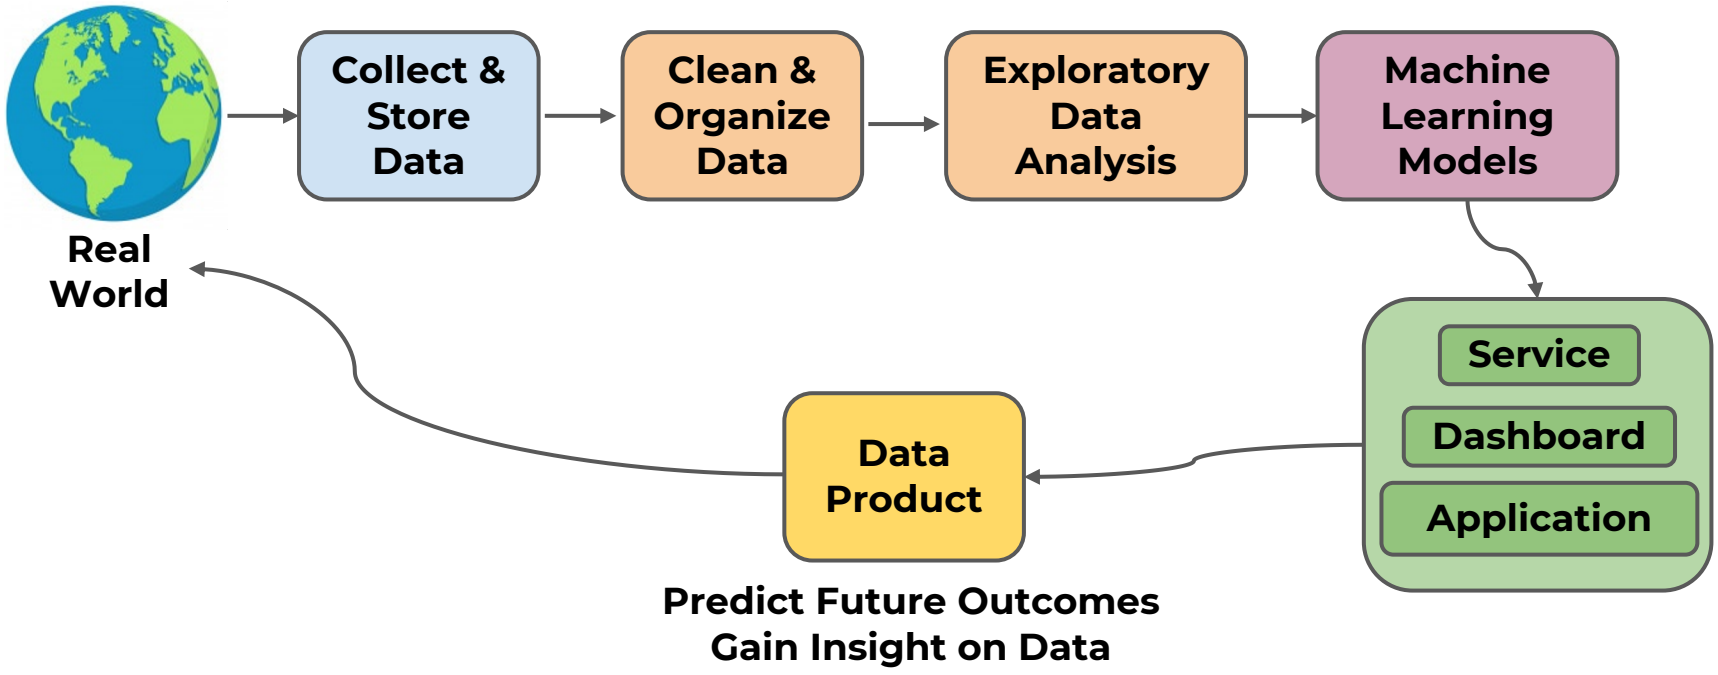
Real World
Collect & Store Data
Clean & Organize Data
Exploratory Data Analysis
Machine Learning Models
Service
Dashboard
Application
Data Product
Predict Future Outcomes
Gain Insight on Data