Intro to Artificial Intelligence



Is it possible
to have
Life
without
Intelligence?

Is it possible to have Intelligence without Life?

🧠 What is AI?

AI as a Discipline

AI is a computing discipline — like Math is a foundational discipline — focused on creating systems that can perceive, reason, and learn. It encompasses core principles, methods, and applications for building intelligent machines.

How Others Define AI

John McCarthy (1956)

"The science and engineering of making intelligent machines."

Father of AI, coined the term at the Dartmouth Conference

Stuart Russell & Peter Norvig

"The study of agents that receive percepts from the environment and perform actions."

Authors of "Artificial Intelligence: A Modern Approach" (standard AI textbook)

Marvin Minsky (MIT AI Lab)

"The science of making machines do things that would require intelligence if done by humans."

Pioneer in AI and cognitive science

Andrew Ng (Stanford/Google Brain)

"The science of making machines smart."

Simplified definition for broad accessibility

IBM

"AI leverages computers and machines to mimic the problem-solving and decision-making capabilities of the human mind."

Industry perspective focusing on practical capabilities

Stanford HAI (Human-Centered AI Institute)

"A field of computer science dedicated to solving cognitive problems commonly associated with human intelligence, such as learning, problem solving, and pattern recognition."

Emphasizes cognitive capabilities and problem-solving

When did AI start?

- 1940's – 1950's
  - First electronic computer developed, and AI is born
- 1960's – 1970's
  - AI research focused on expert systems
  - First chatbot, ELIZA, is created in 1966 and machine learning emerges
- 1980's – 1990's
  - AI research shifts to focus on neural networks and robotics emerges
- 2000's – present
  - Focus on use of deep learning and big data
  - Ethical issues raised





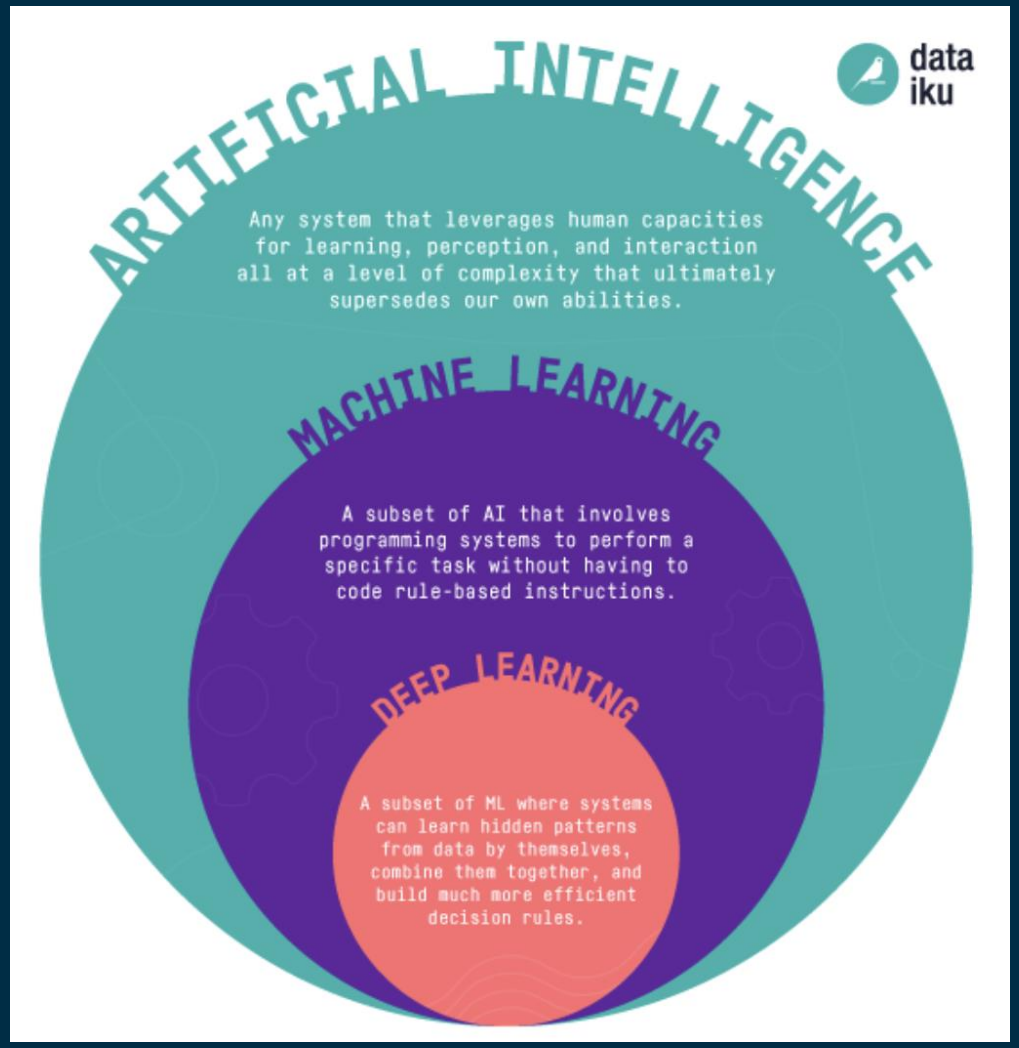
1. Storage and Processing Power of Computers

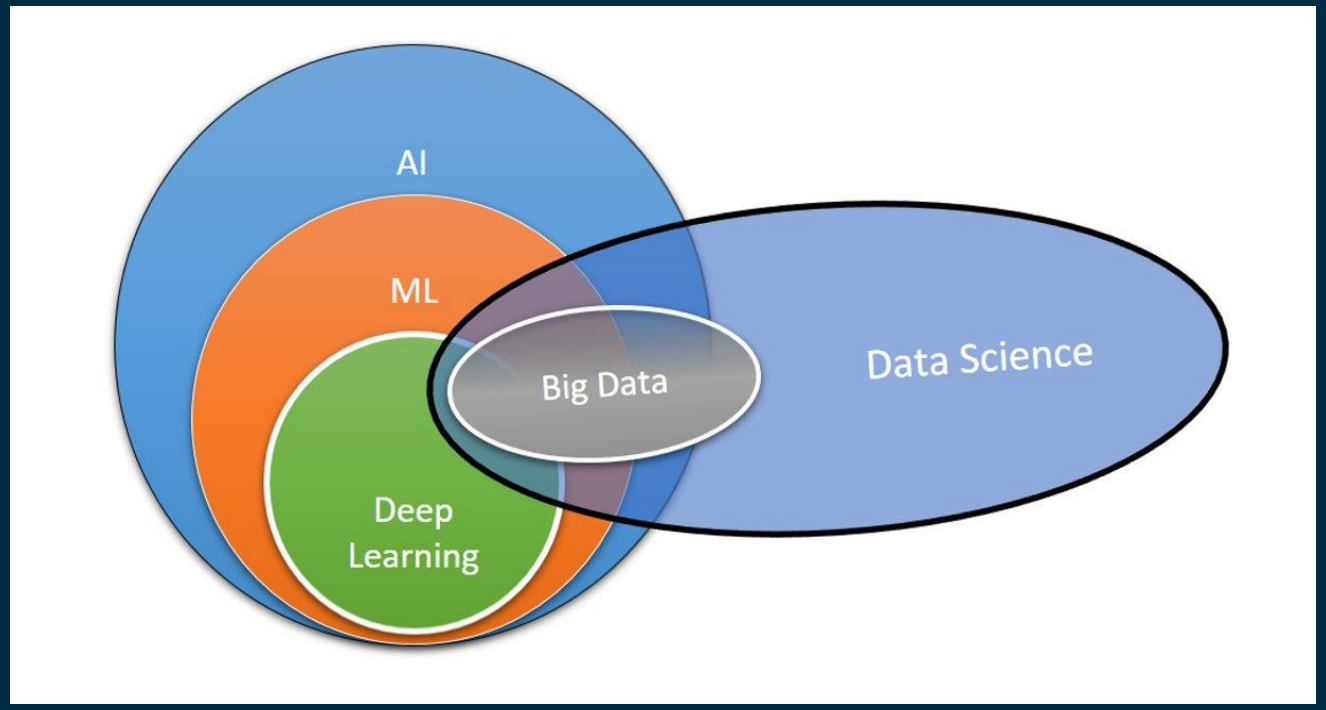
2. Vast Amounts of Data

TWO Contributing Factors

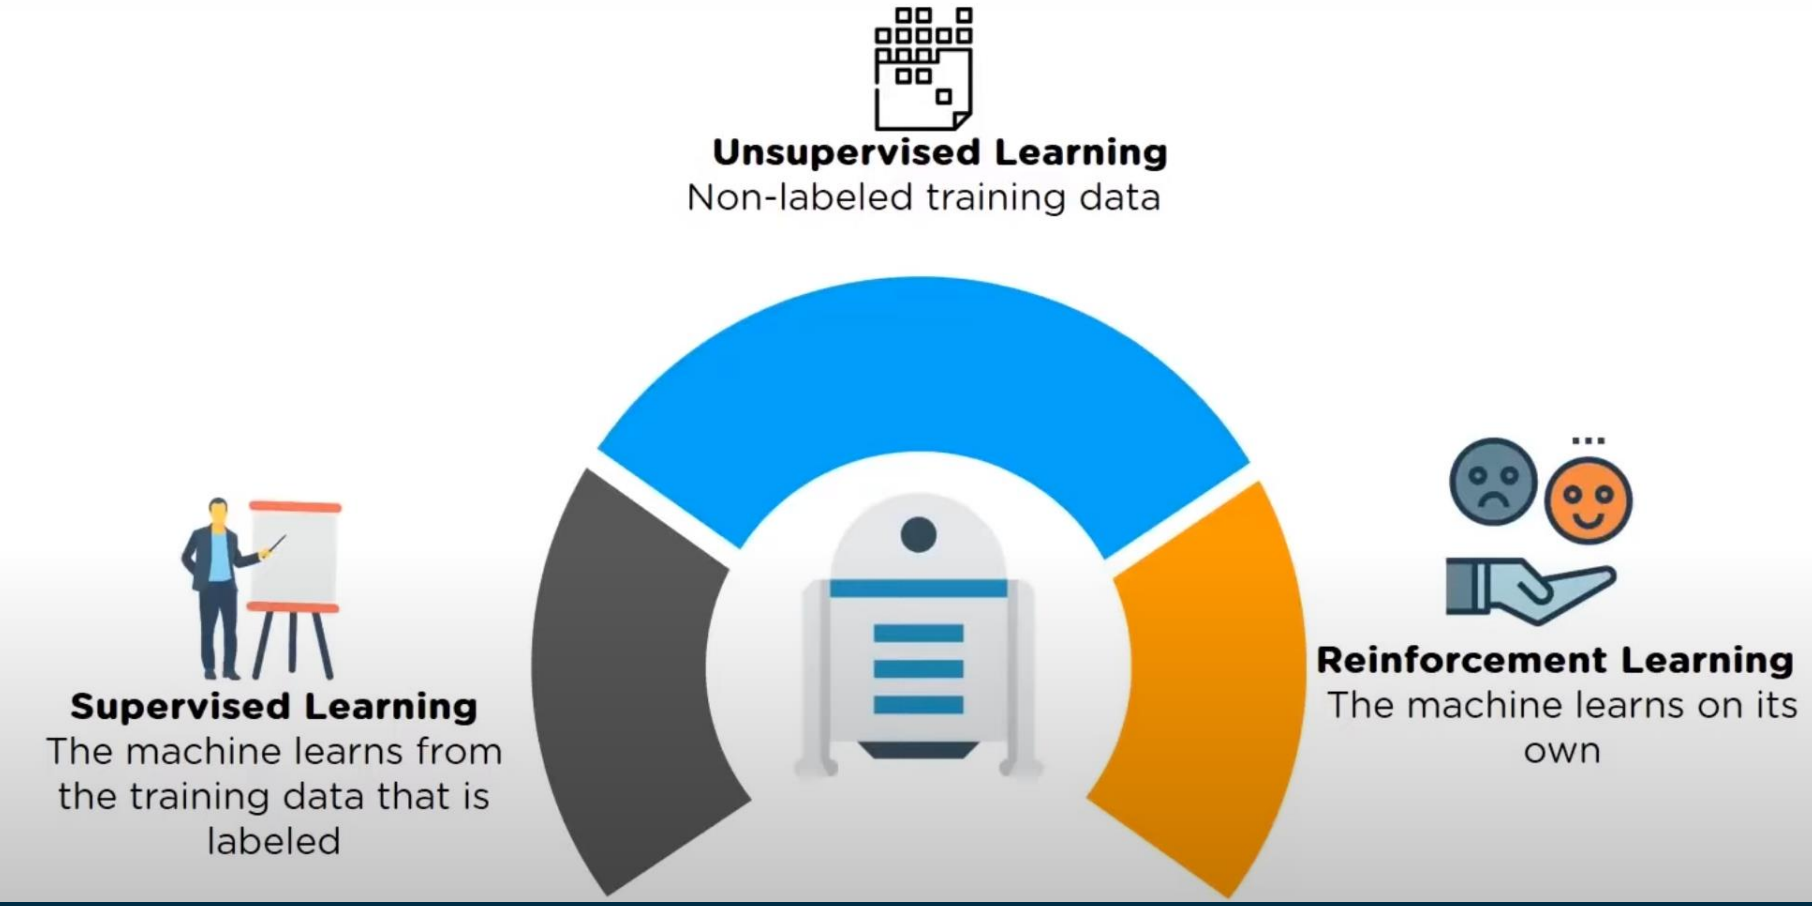


AI vs. ML vs. Deep Learning





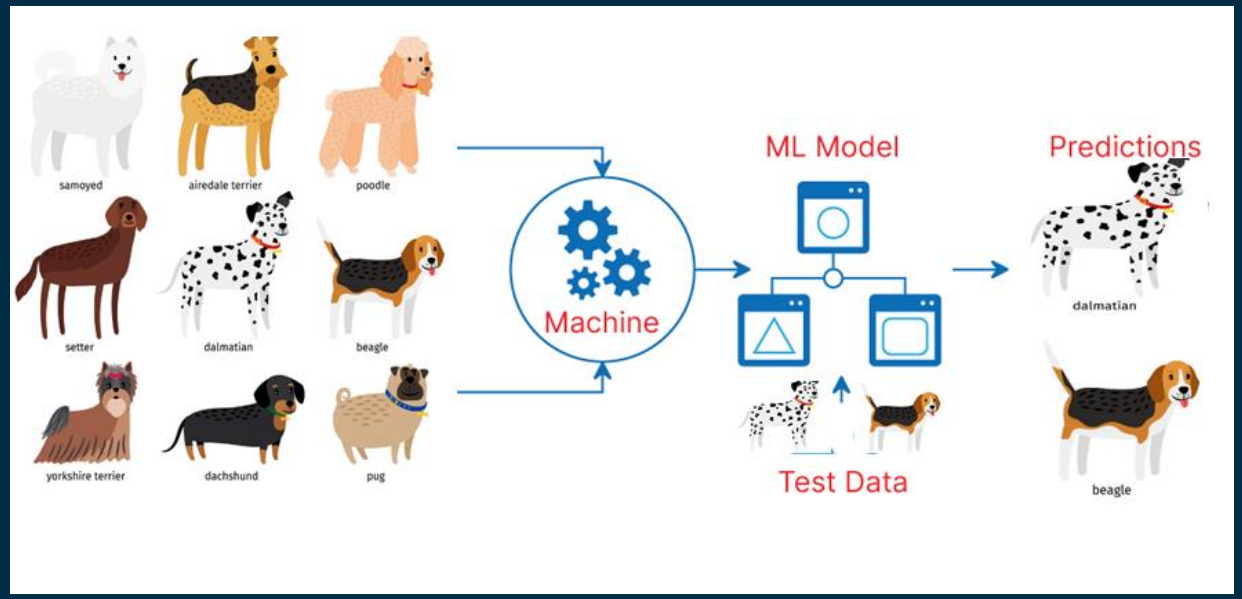
Types of Machine Learning

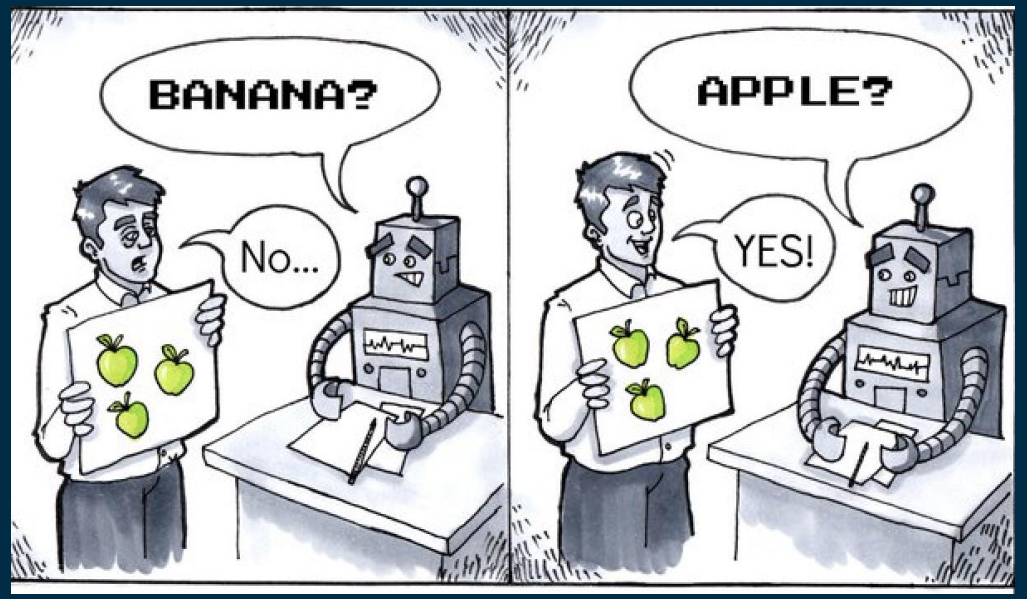


Source: SimpliLearn https://www.youtube.com/watch?v=olFxW7kdtP8

Supervised Learning

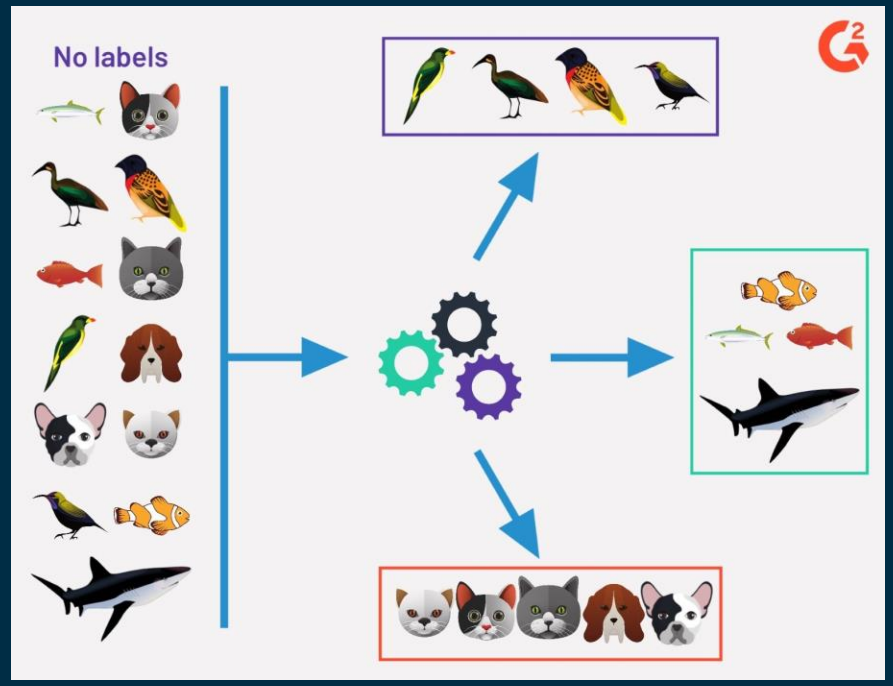
- AI receives labeled data
- Tries to make predictions about new data

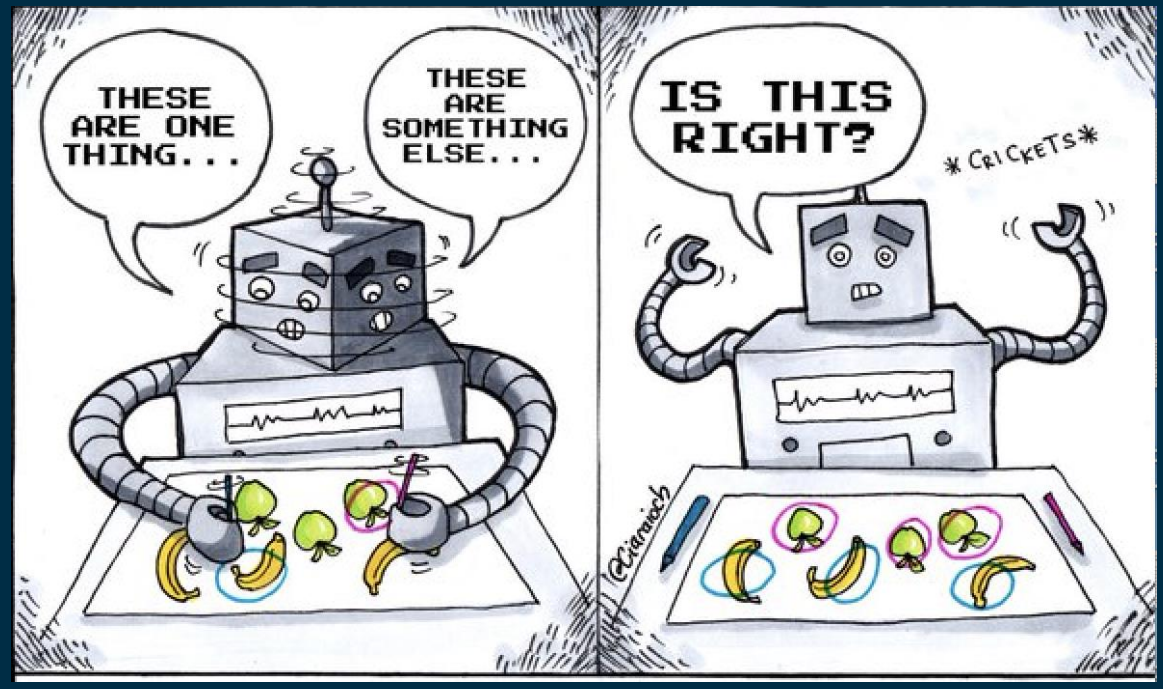




Unsupervised Learning

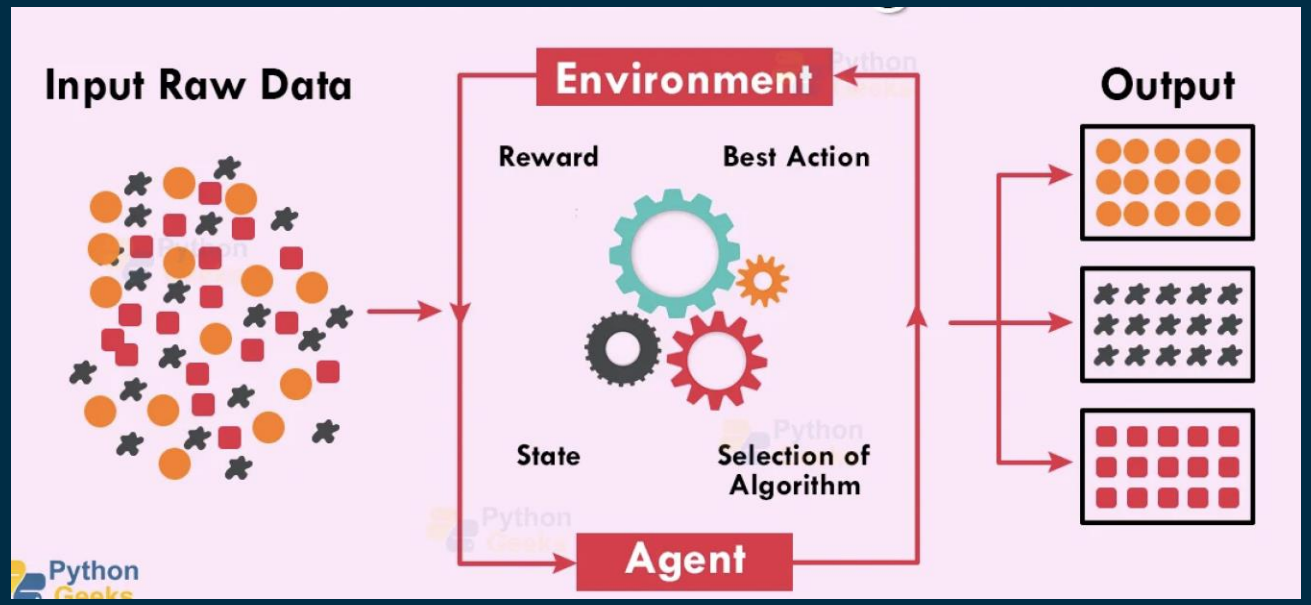
- AI receives unlabeled data
- Identifies patterns to make predictions about new data





Reinforcement Learning

- AI's actions are rewarded or punished to incentivize desired behavior

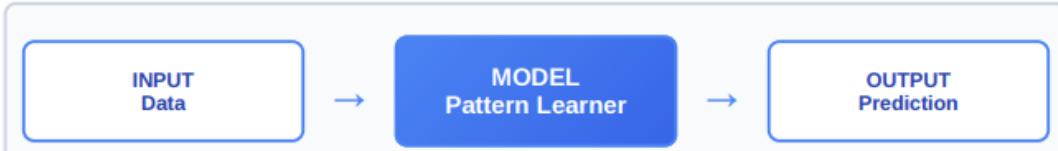




🤖 What is a Model?

Understanding the foundation of machine learning

The Basic Concept



A model is a learned pattern from data that makes predictions or decisions.

Familiar Examples (You Already Know Models!)

🌤️ Weather Forecast
Past weather data → Model → Tomorrow's temperature

🗺️ GPS Navigation
Roads & traffic → Model → Best route home

📧 Spam Filter
Email content → Model → Spam or Not Spam

🏥 Medical Diagnosis
Symptoms & tests → Model → Likely disease

✓ A Model IS...

- ✓ A pattern learned from data
- ✓ A simplified representation of reality
- ✓ A tool for making predictions
- ✓ Only as good as its training data
- ✓ Constantly improving with more data

✗ A Model is NOT...

- ✗ Magic or "true AI" that thinks
- ✗ Perfect (models make mistakes)
- ✗ A replacement for human judgment
- ✗ One-size-fits-all solution
- ✗ Useful without good training data

💡 Key Insight: Models learn from examples, just like humans do!
The more quality examples, the better the model performs.

🤖 Machine Learning Models: A Practical Guide

Choose the right approach based on your data and goal

📊 Supervised Learning

> ✓ You have labeled data (inputs + correct answers)

Classification

Predicting categories or labels

- Is this email spam or not spam?
- Is this image a cat or dog?
- Does patient have disease?
- Will customer churn?

Regression

Predicting numerical values

- What will the house price be?
- How many units will sell?
- What temperature tomorrow?
- Student's test score?

Common algorithms: Decision Trees, Random Forest, Neural Networks, SVM, Linear/Logistic Regression, Naive Bayes

🔍 Unsupervised Learning

> ✓ You have unlabeled data (no "right answers")

Clustering

Group similar items together

- Customer segmentation
- Document organization
- Image compression
- Anomaly detection

Dimensionality Reduction

Simplify complex data

- Data visualization
- Feature extraction
- Noise reduction
- Speed up algorithms

Association Rules

Find patterns and relationships

- Market basket analysis
- Product recommendations
- "Frequently bought together"

Common algorithms: K-Means, DBSCAN, Hierarchical Clustering, PCA, t-SNE, Apriori, FP-Growth

🤖 Machine Learning Models: A Practical Guide

Choose the right approach based on your data and goal

⚖️ Semi-Supervised Learning

> ✓ You have SOME labeled + LOTS of unlabeled data

Hybrid Classification

Use limited labels efficiently

- Medical imaging (few labels)
- Text classification
- Speech recognition
- Web page categorization

Why use it: Labeling is expensive/time-consuming. Leverage unlabeled data to improve performance!

Common algorithms: Self-training, Co-training, Semi-supervised SVM, Graph-based methods

🎮 Reinforcement Learning

> ✓ Agent learns through trial & error (rewards/penalties)

Sequential Decision Making

Learn optimal strategy over time

- Game playing (Chess, Go, Atari)
- Robot navigation & control
- Self-driving cars
- Resource management
- Trading strategies

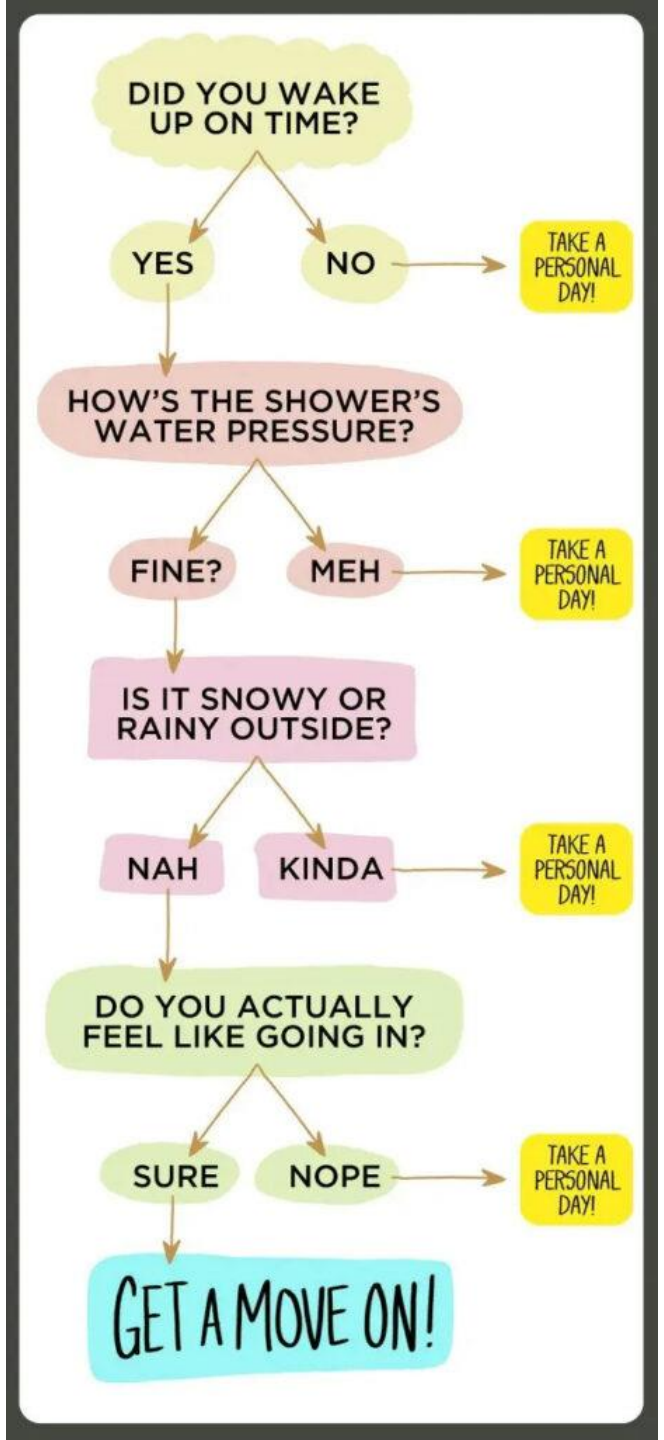
Key concept: Agent takes action → Environment gives reward/penalty → Agent learns better policy

Common algorithms: Q-Learning, SARSA, Deep Q-Networks (DQN), Policy Gradients, Actor-Critic

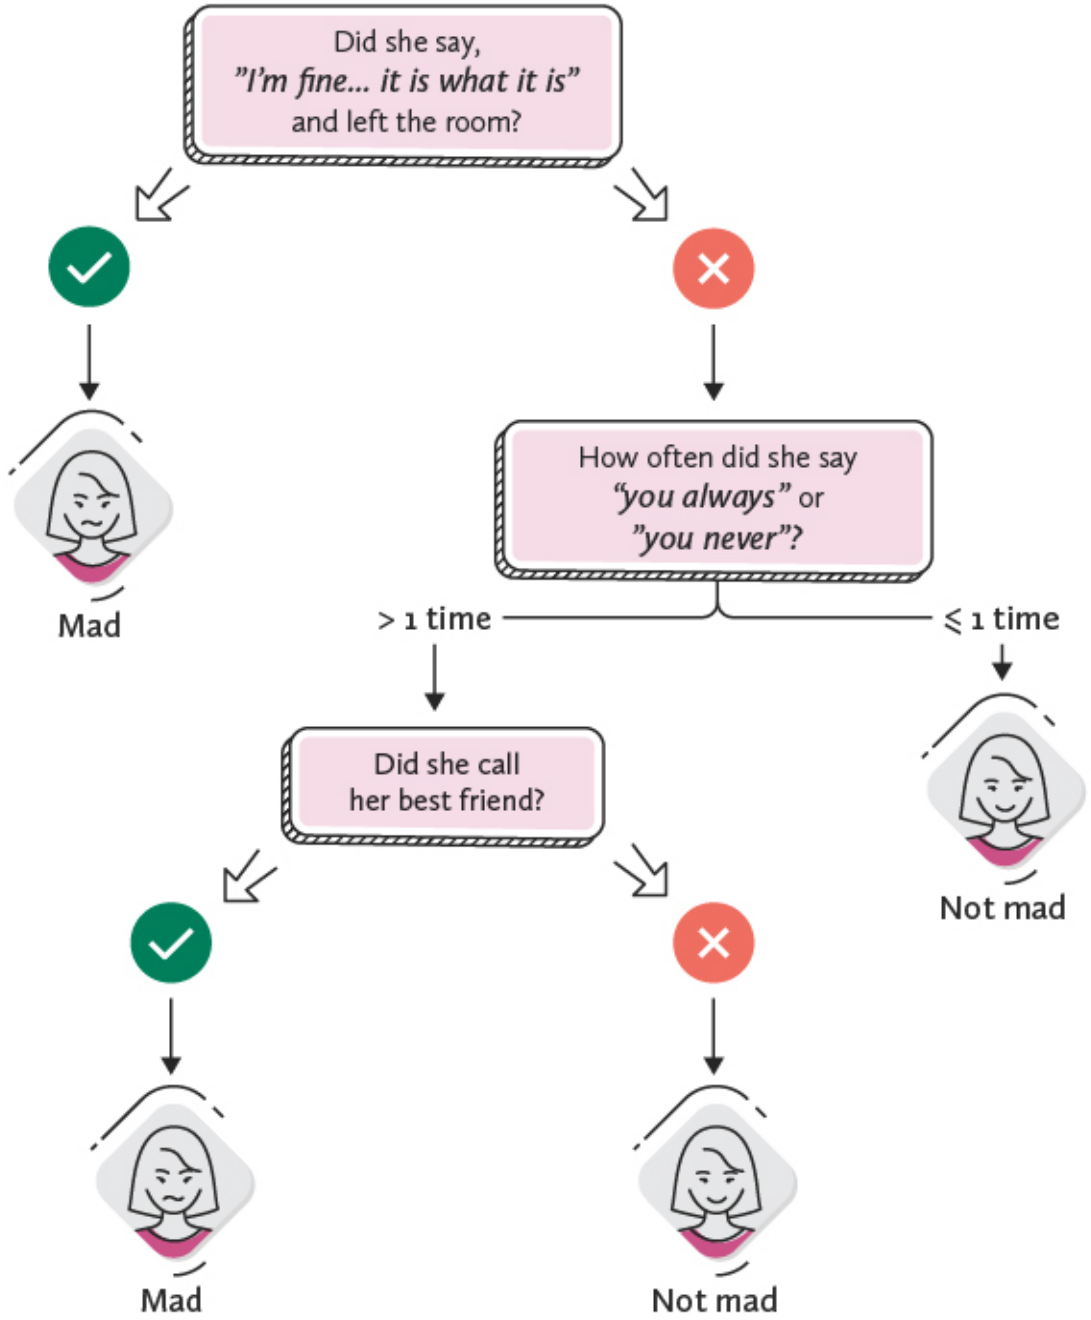
💡 Start by asking: "What kind of data do I have?" and "What am I trying to achieve?"

Unplugged AI Activity

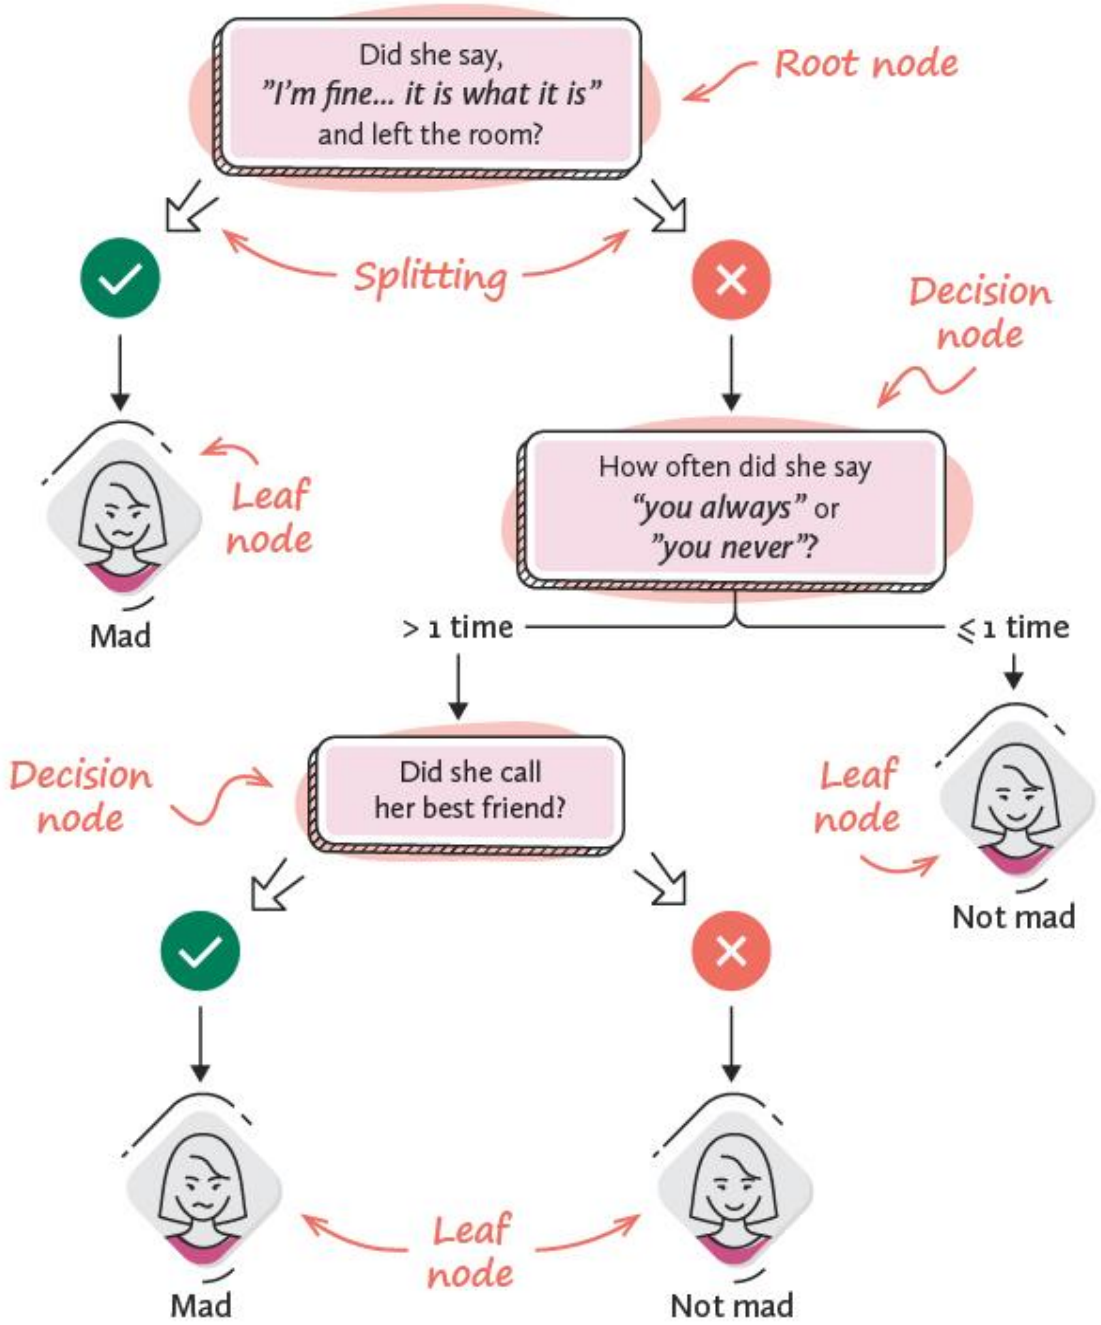
Classification Using a Decision Tree



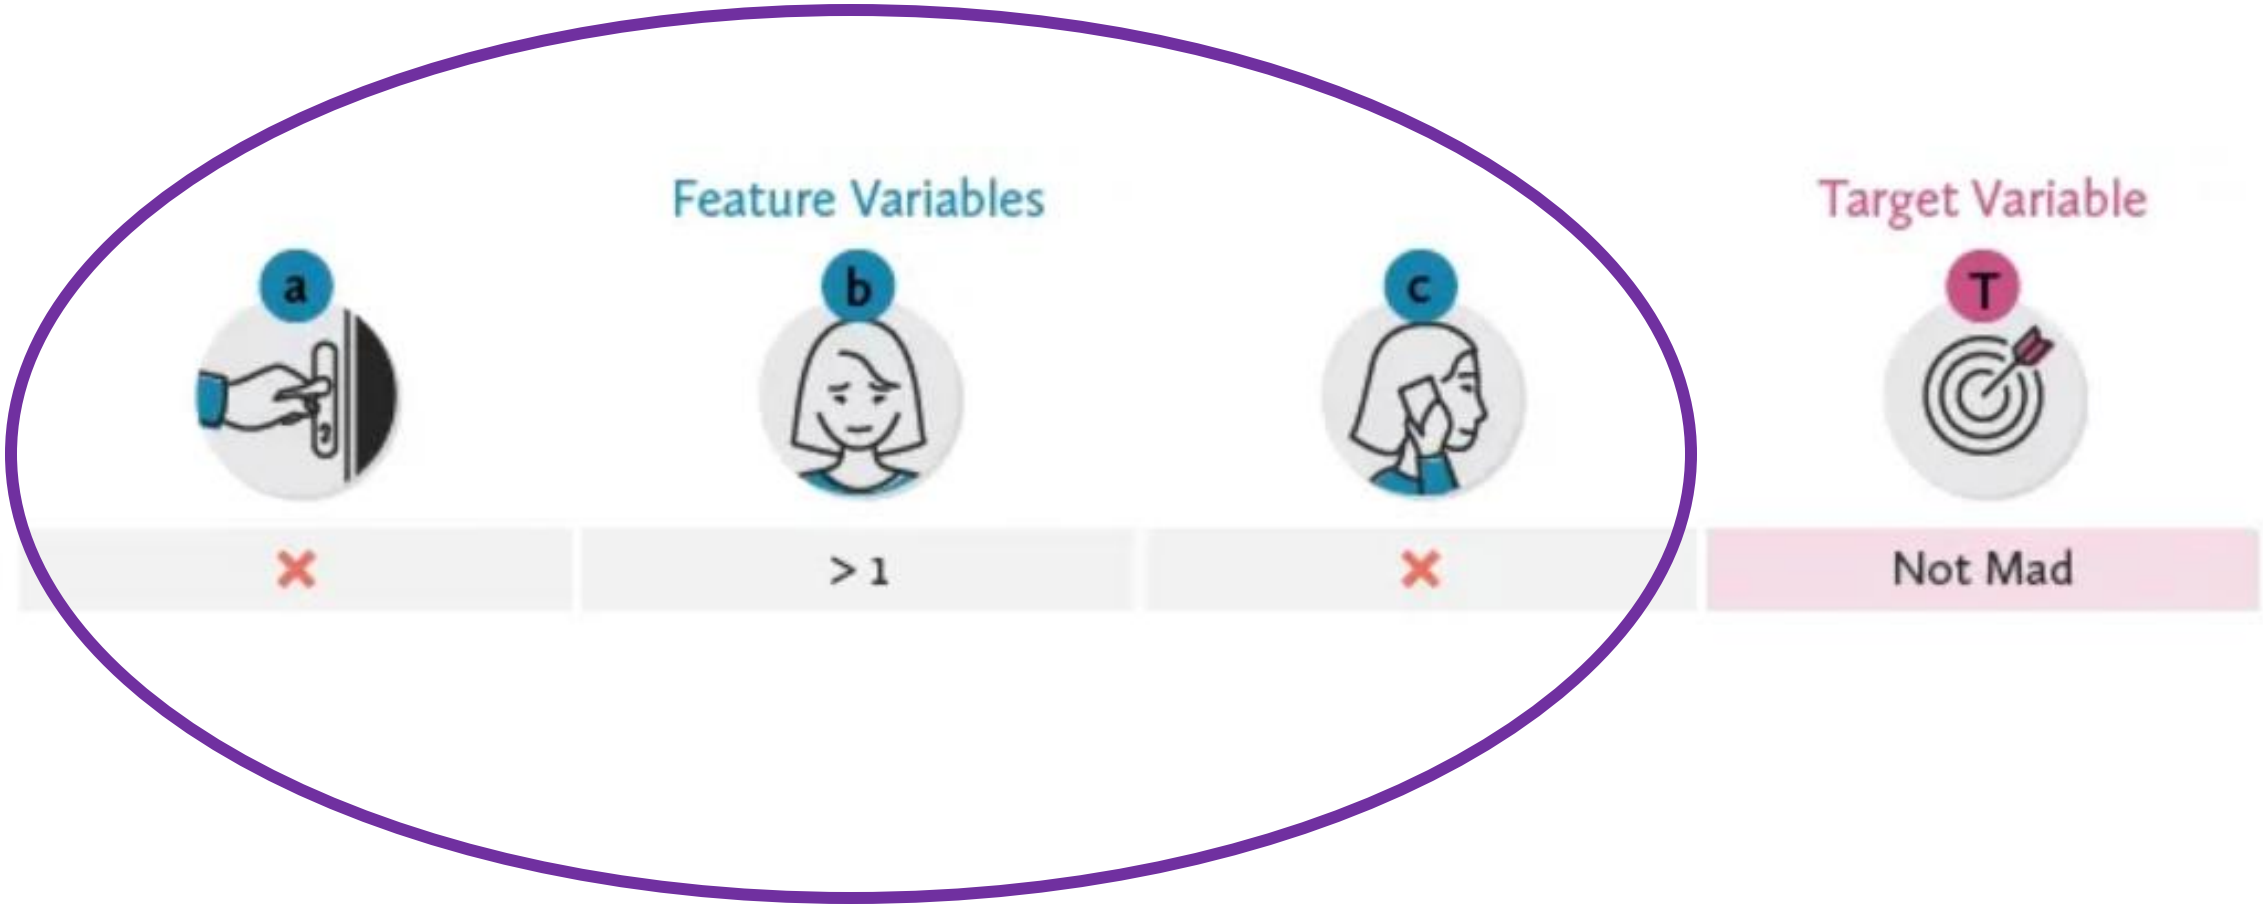
Classification Using a Decision Tree



Decision Nodes & Leaf Nodes



What is Information Gain?
What Question will you ask FIRST?



Unplugged AI Activity

Classification Using a Decision Tree

You will create your own **MODEL** using a Decision Tree

- AI classifies data based on patterns
- Classification models are not perfect

Unplugged AI Activity

The Backstory!

- The monkeys in the zoo are cute, BUT they BITE!
- We already know if our monkeys bite – we have a to determine this.
- We are getting NEW monkeys visiting us.
- We need a way to find out if they BITE!

Unplugged AI Activity

There are 2 Datasets

- **Training – Train the Model**
  - **Think of RULES to distinguish the ones that BITE!**
  - **Some features to look for:**
    - **Shape of mouth**
    - **Teeth showing**
    - **Eye open**
    - **Wearing accessories**

- **Testing – test the model you created**

Categorizing the Training Data

Biting

1, 2, 5, 9, 10, 14, 15, 16, 17, 28, 33, 35, 36

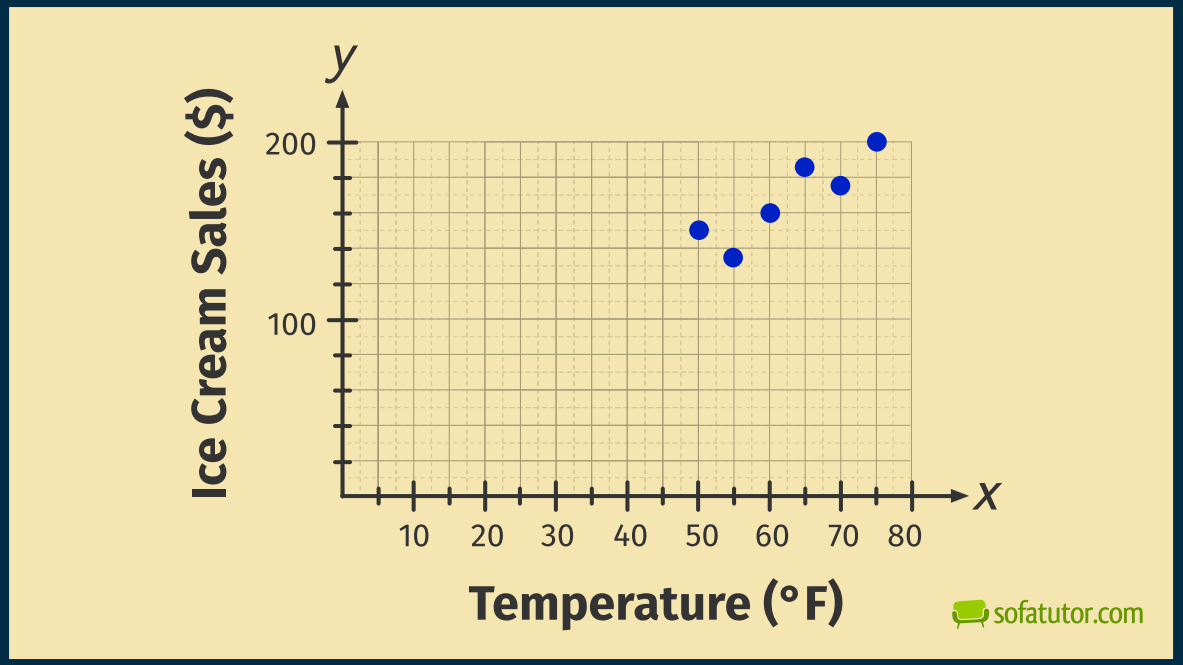
Non-Biting

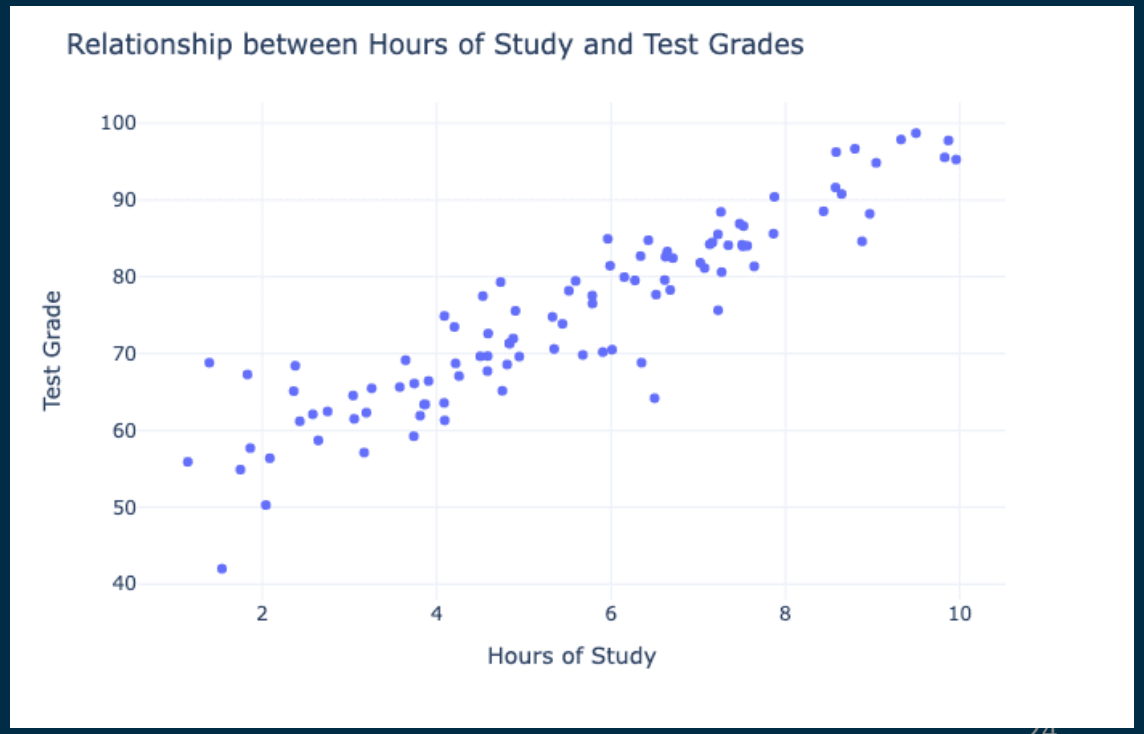
4, 7, 12, 19, 22, 23, 24, 25, 30, 32, 37, 38, 39, 40

Categorizing the Training Data

Monkey#	Mask	Spiky Hair	2EyesOpen	1EyeOpen	EyesClosed	TeethShown	HairBow	Band-Aid	Glasses	MouthOpen	XEyes	TongueOut	Biting
1													
2													
4													
5													
7													
9													
10													
12													
14													
15													
16													
17													
19													
22													
23													
24													
25													
28													
30													
32													
33													
35													
36													
37													
38													
39													
40													

What is a Model?

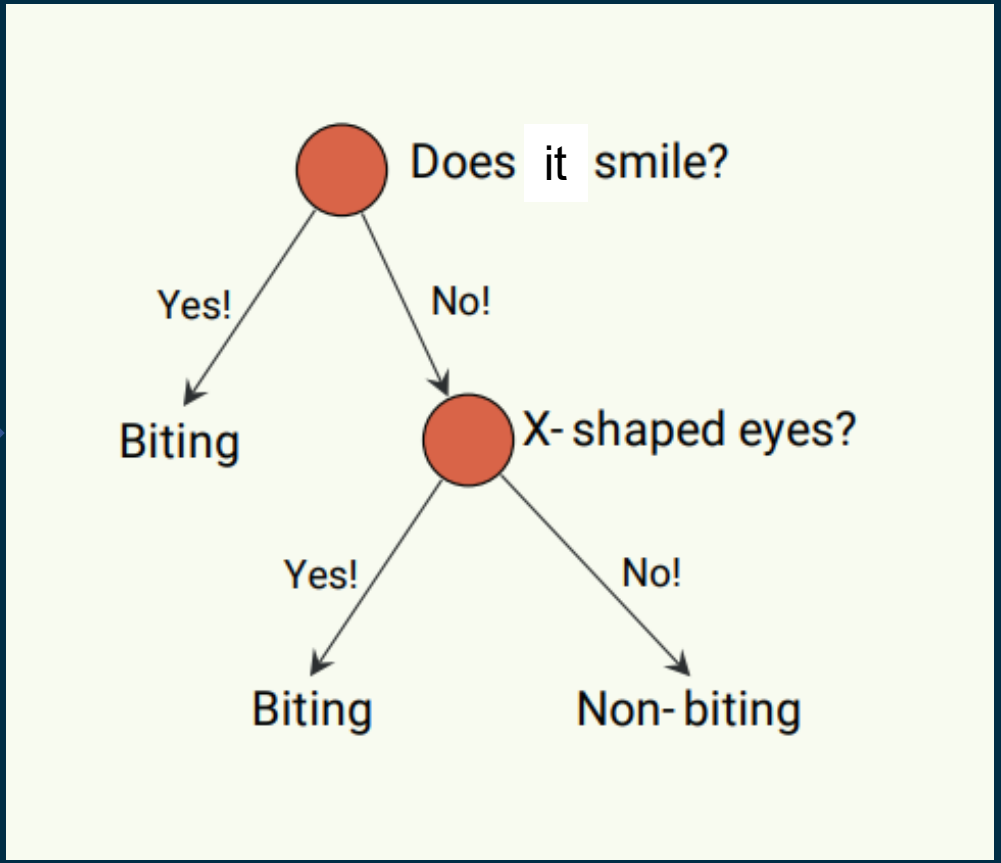


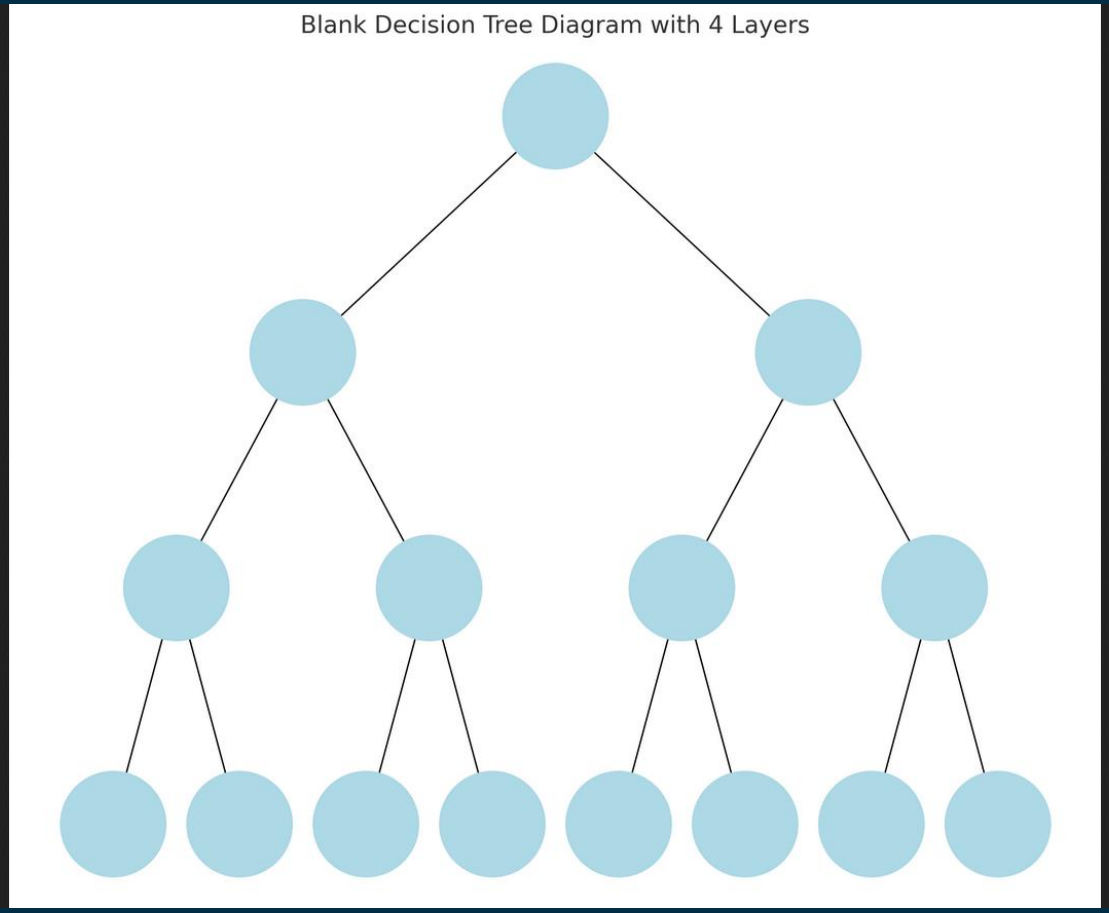


Unplugged AI Activity

Create a Decision Tree to Categorize the Monkeys into Biting and Non-Biting

Example Decision Tree




Blank Decision Tree Diagram with 4 Layers

Unplugged AI Activity

- Open your Test Data
- For Each Image:
  - Categorize the Monkeys into Biting and Non-Biting

Unplugged AI Activity

STOP

Unplugged AI Activity

Which Monkeys BITE?

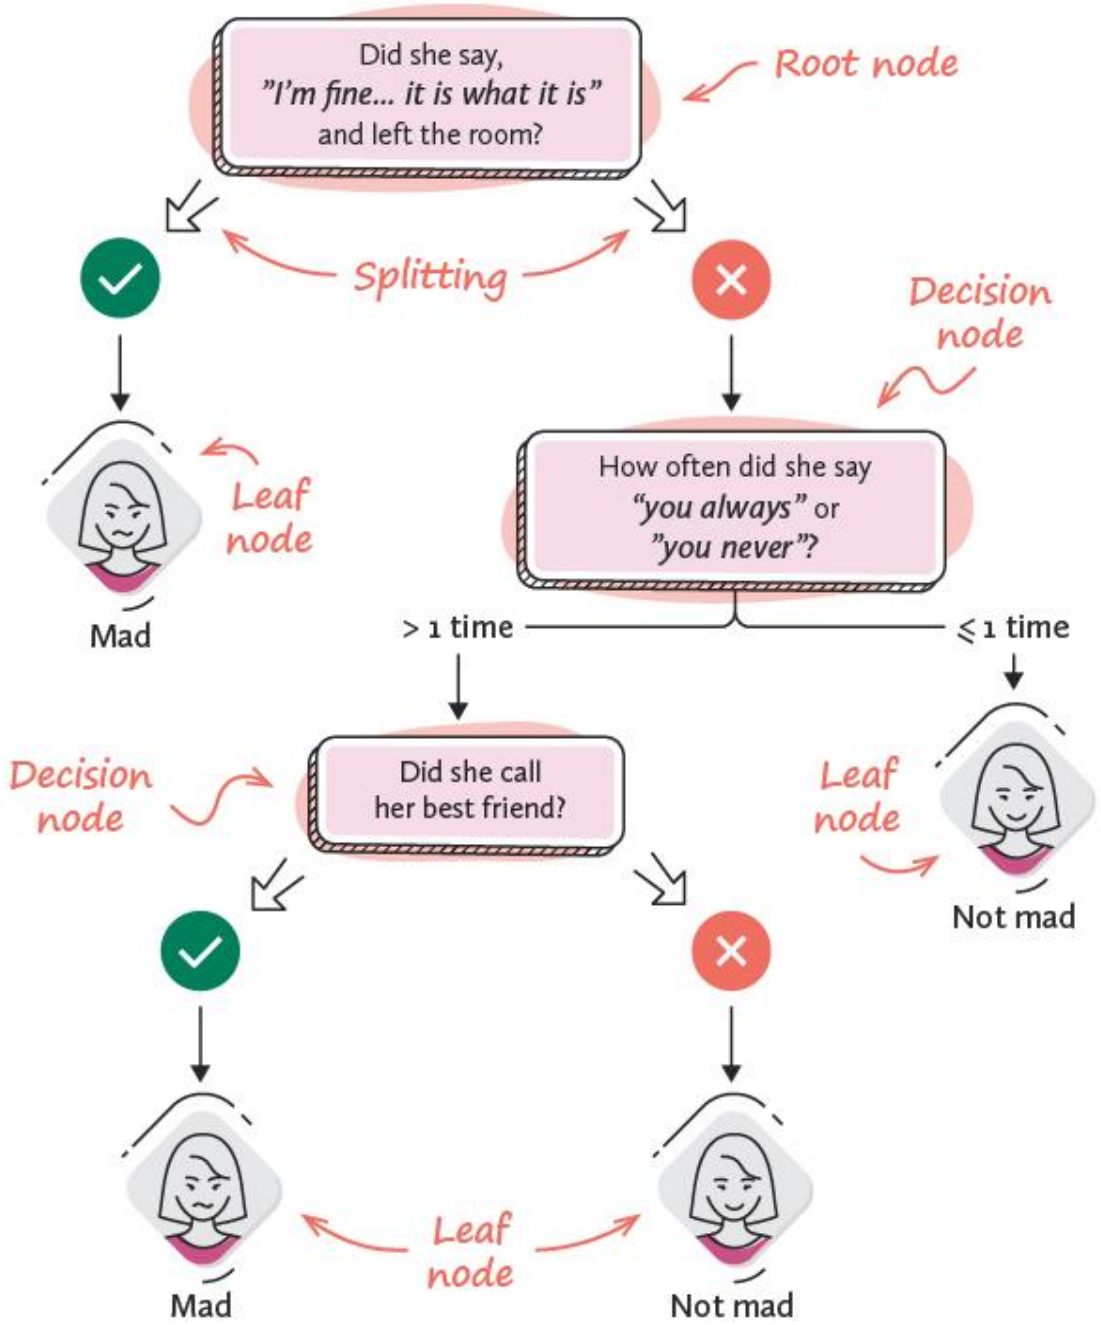
biting: 6, 13, 18, 34
non-biting; 3, 8, 11, 20, 21, 26, 27, 29, 31

- Who had the **MOST** successful model?
- Were there any possible issues?

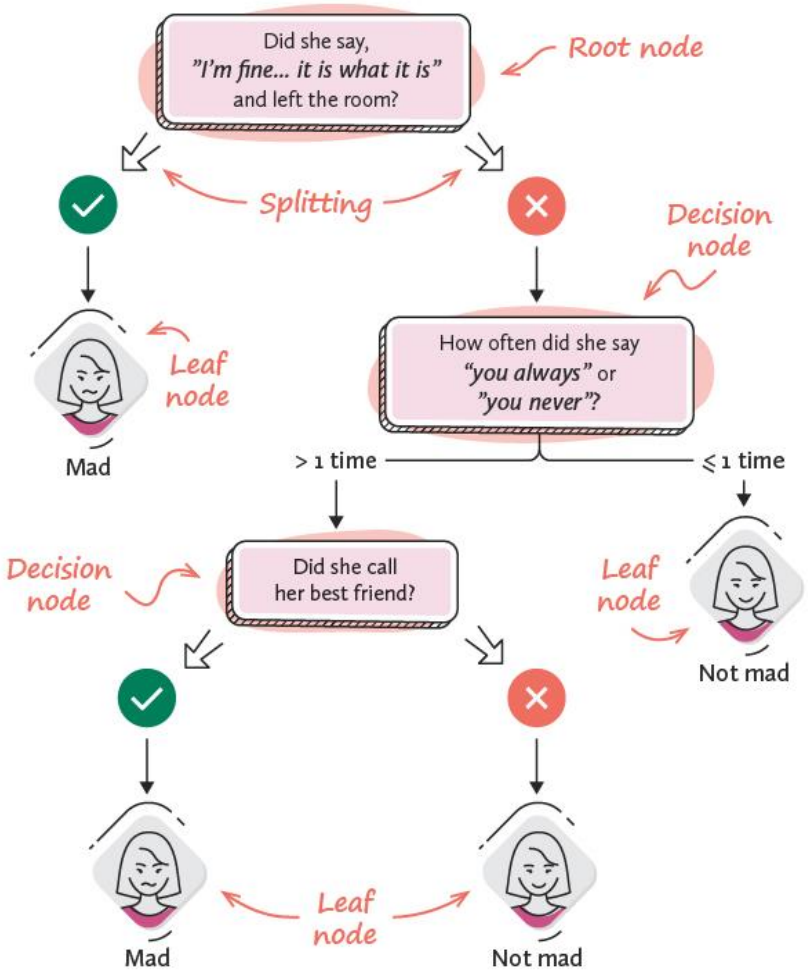
Unplugged AI Activity

Reflection!

Decision Nodes & Leaf Nodes



A Decision Tree Model for this Situation

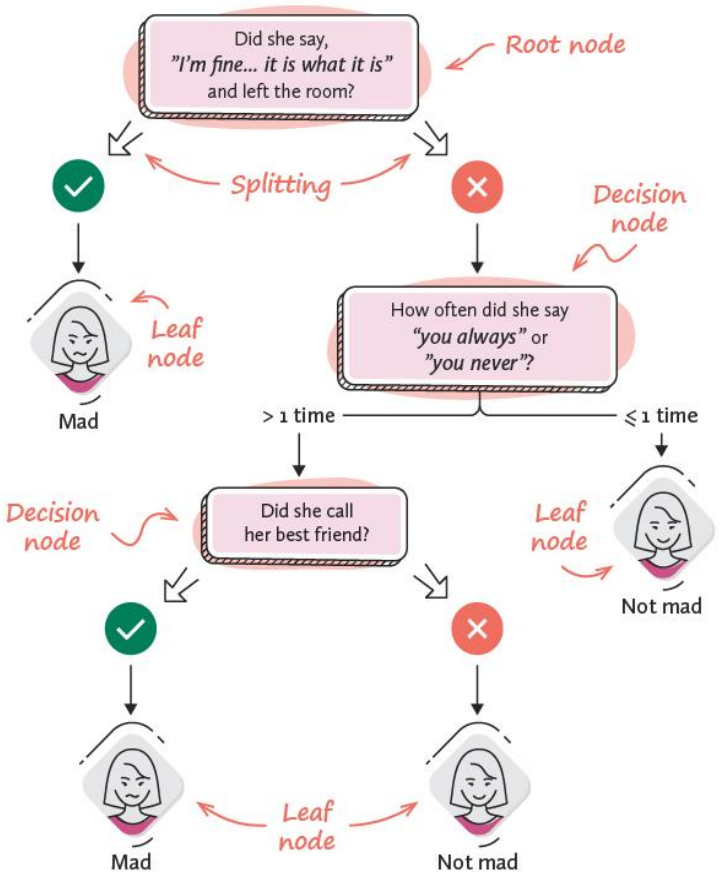


The Actual Data Collected By You!

Row	Said "I'm fine"	"Always/never " freq	Called friend	Actual: Mad
1	Yes	>1	Yes	Mad
2	Yes	>1	No	Mad
3	Yes	≤1	Yes	Not mad
4	Yes	≤1	No	Mad
5	No	>1	Yes	Not mad
6	No	>1	No	Not mad
7	No	≤1	Yes	Mad
8	No	≤1	No	Not mad

Data collected from your last 8 experiences – when you had to figure out if your girlfriend was mad … or not!

How accurate was the model?



Row	Said "I'm fine"	"Always/never" freq	Called friend	Actual: Mad	Model Predicts	Result
1	Yes	>1	Yes	Mad	Mad	✓ Correct
2	Yes	>1	No	Mad	Mad	✓ Correct
3	Yes	≤1	Yes	Not mad	Mad	✗ False Positive
4	Yes	≤1	No	Mad	Mad	✓ Correct
5	No	>1	Yes	Not mad	Mad	✗ False Positive
6	No	>1	No	Not mad	Not mad	✓ Correct
7	No	≤1	Yes	Mad	Not mad	✗ False Negative
8	No	≤1	No	Not mad	Not mad	✓ Correct

Do you understand these?

True Positive (TP), False Positive (FP), True Negative (TN), False Negative (FN)

How do we EVALUATE the model?

Row	Said "I'm fine"	"Always/ never" freq	Called friend	Actual: Mad	Model Predicts	Result
1	Yes	>1	Yes	Mad	Mad	✓ Correct
2	Yes	>1	No	Mad	Mad	✓ Correct
3	Yes	≤1	Yes	Not mad	Mad	✗ False Positive
4	Yes	≤1	No	Mad	Mad	✓ Correct
5	No	>1	Yes	Not mad	Mad	✗ False Positive
6	No	>1	No	Not mad	Not mad	✓ Correct
7	No	≤1	Yes	Mad	Not mad	✗ False Negative
8	No	≤1	No	Not mad	Not mad	✓ Correct

Confusion Matrix

	Predicted: Mad	Predicted: Not Mad
Actual: Mad	3 (True Positive)	1 (False Negative)
Actual: Not Mad	2 (False Positive)	2 (True Negative)

- **True Positives (TP)**: 3 (rows 1, 2, 4)
- **False Positives (FP)**: 2 (rows 3, 5)
- **False Negatives (FN)**: 1 (row 7)
- **True Negatives (TN)**: 2 (rows 6, 8)

How do we EVALUATE the model?

Row	Said "I'm fine"	"Always/never" freq	Called friend	Actual: Mad	Model Predicts	Result
1	Yes	>1	Yes	Mad	Mad	✓ Correct
2	Yes	>1	No	Mad	Mad	✓ Correct
3	Yes	≤1	Yes	Not mad	Mad	✗ False Positive
4	Yes	≤1	No	Mad	Mad	✓ Correct
5	No	>1	Yes	Not mad	Mad	✗ False Positive
6	No	>1	No	Not mad	Not mad	✓ Correct
7	No	≤1	Yes	Mad	Not mad	✗ False Negative
8	No	≤1	No	Not mad	Not mad	✓ Correct

Precision

Precision = TP / (TP + FP)

= 3 / (3 + 2) = 3/5 = 0.60 or 60%

This means when the model predicts "Mad," it's correct 60% of the time.

Recall (Sensitivity)

Recall = TP / (TP + FN)

= 3 / (3 + 1) = 3/4 = 0.75 or 75%

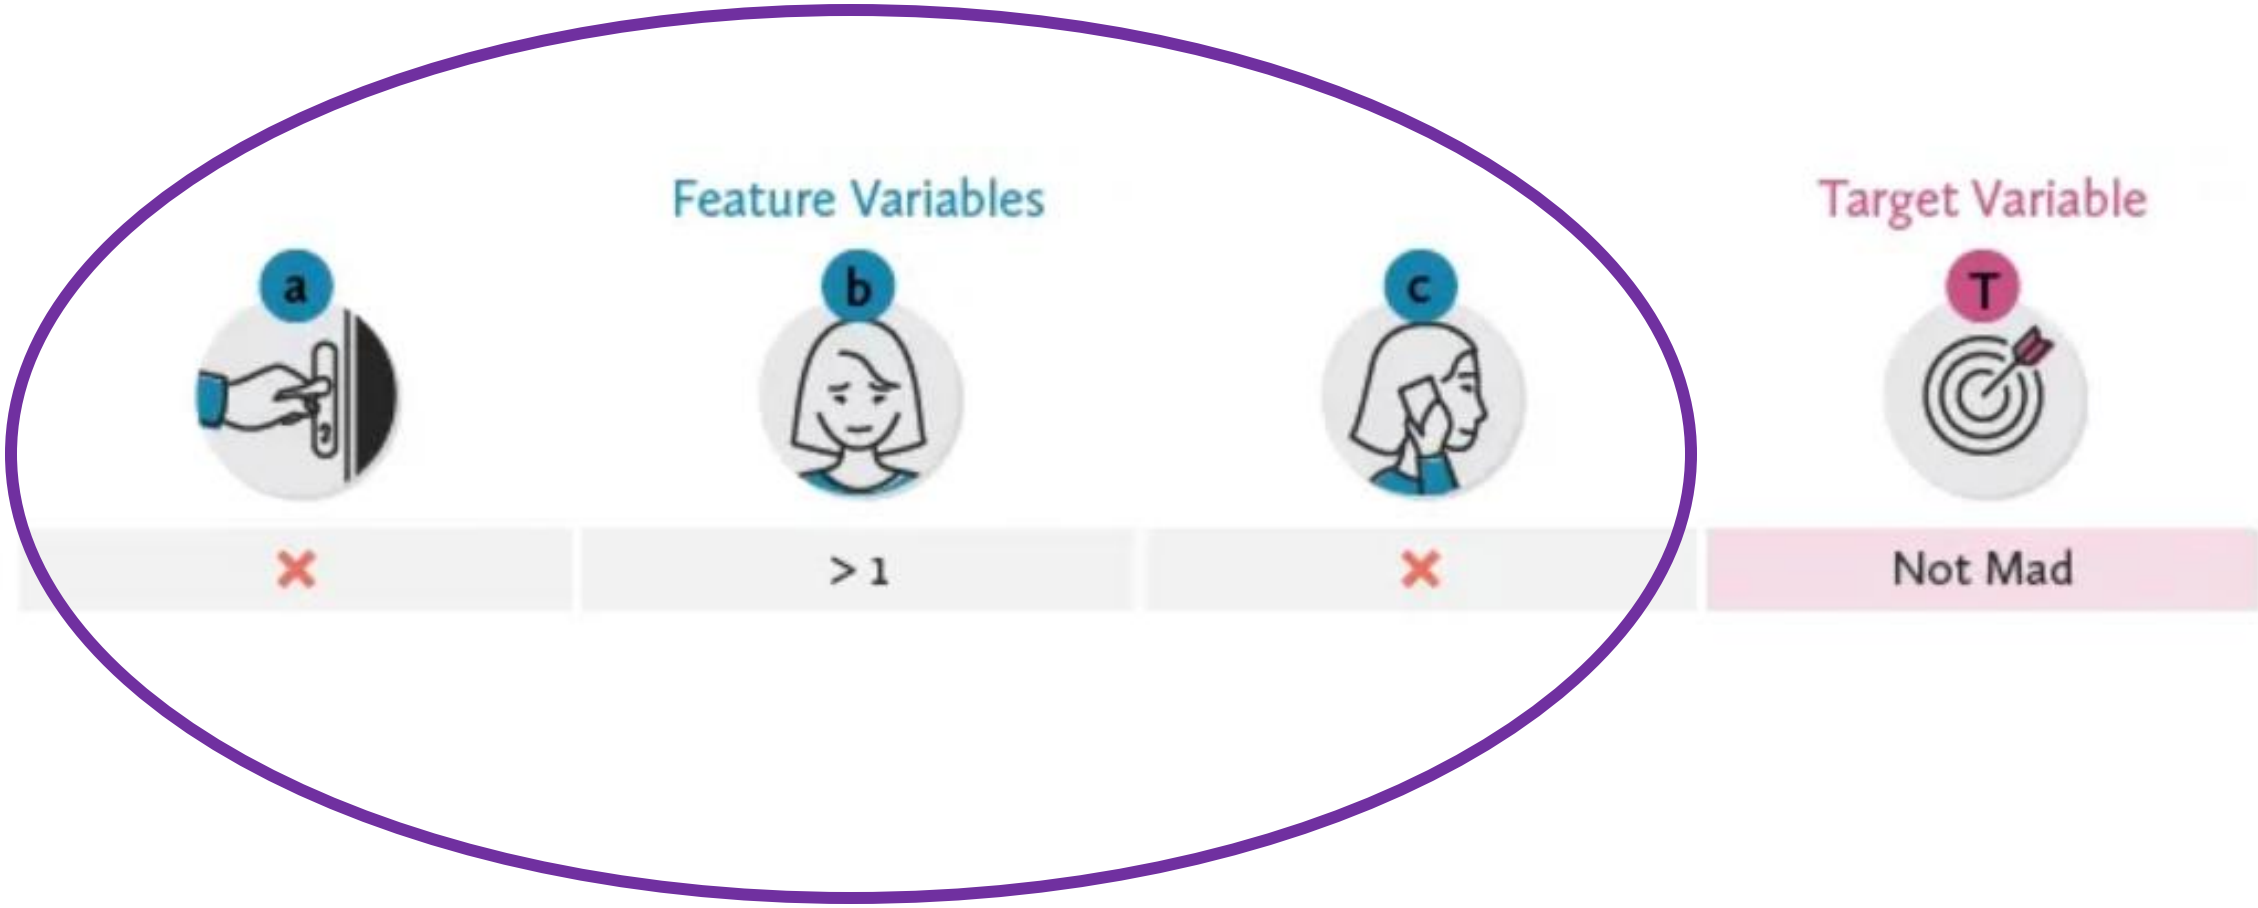
This means the model correctly identifies 75% of all actual "Mad" cases.

Accuracy

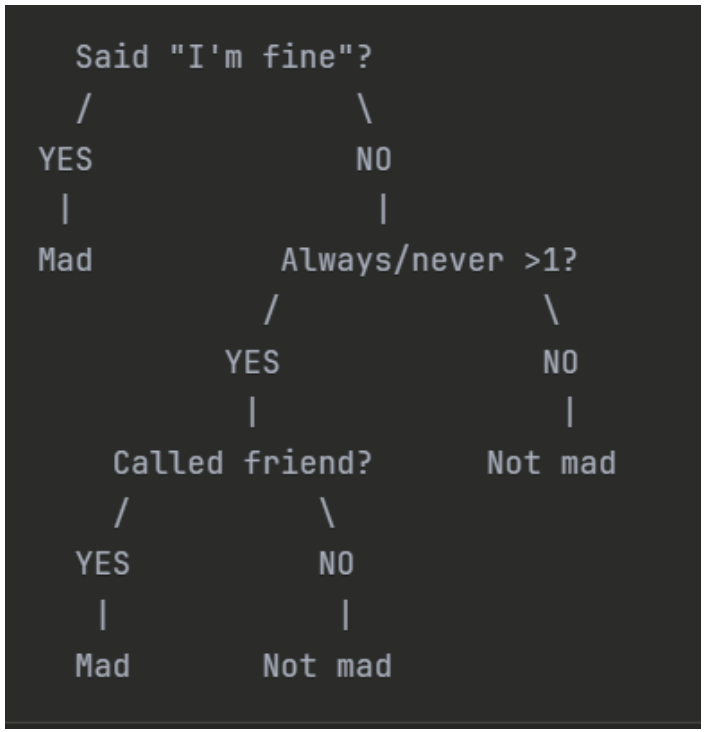
Accuracy = (TP + TN) / Total

= (3 + 2) / 8 = 5/8 = 0.625 or 62.5%

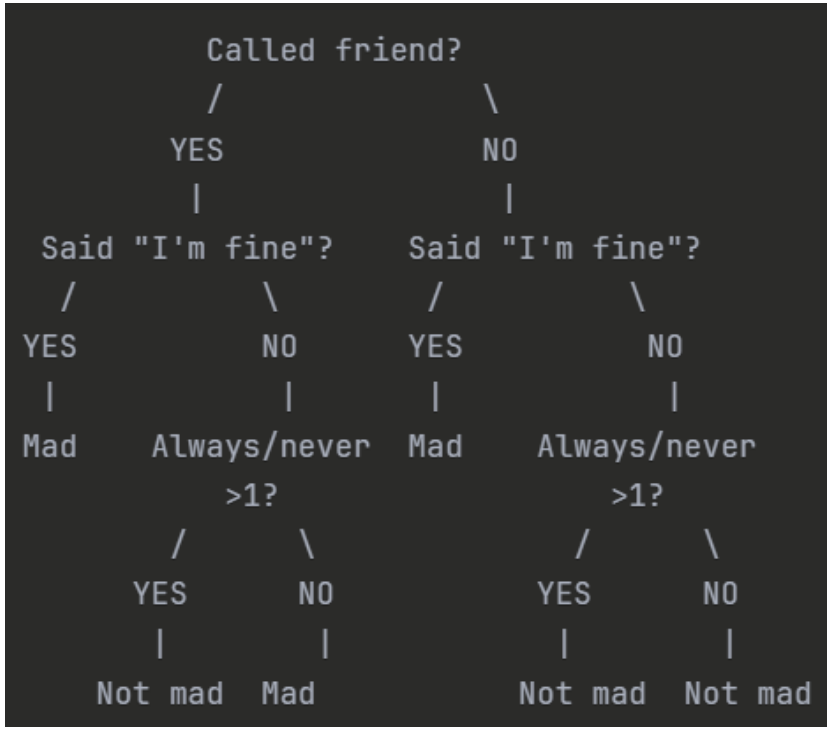
What Question will you ask FIRST?
Information GAIN!



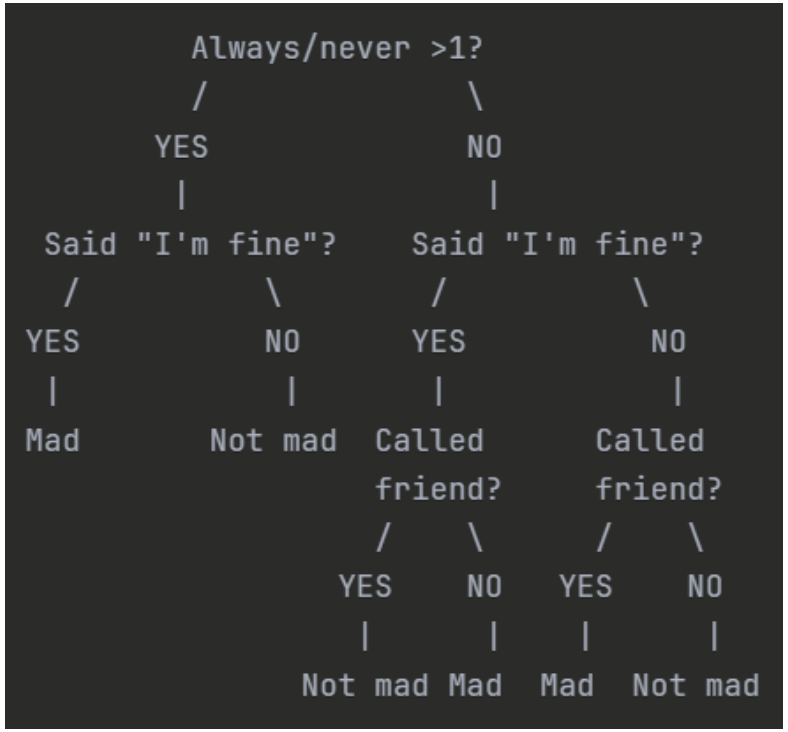
What is Information Gain?
What Question will you ask FIRST?



First Question: Said: "I'm fine"?



First Question: Called a friend?



First Question: always/never >1??

EVALUATE Tree 2 or Tree 3

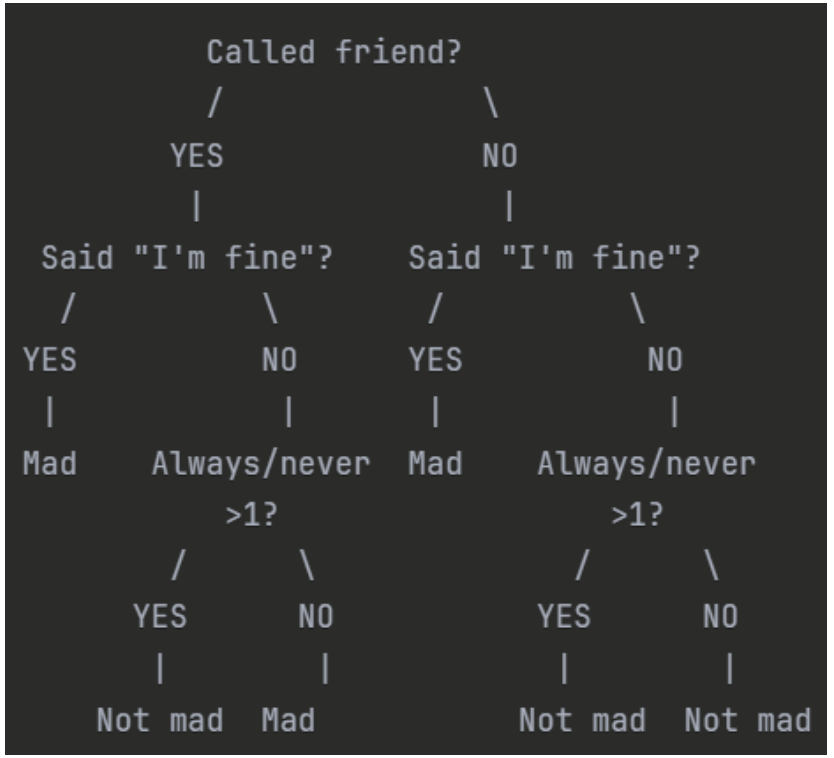
Your Task

1. Pick Tree 2 or Tree 3
2. For each row, trace through the tree and write down the model's prediction
3. Compare each prediction to the actual outcome (TP, TN, FP, or FN)
4. Build the **Confusion Matrix**
5. Calculate **Precision, Recall, and Accuracy**

EVALUATE Tree 2

Row	Said "I'm fine"	"Always/ never" freq	Called friend	Actual : Mad	Model Predicts	Result
1	Yes	>1	Yes	Mad		
2	Yes	>1	No	Mad		
3	Yes	≤1	Yes	Not mad		
4	Yes	≤1	No	Mad		
5	No	>1	Yes	Not mad		
6	No	>1	No	Not mad		
7	No	≤1	Yes	Mad		
8	No	≤1	No	Not mad		

- True Positives (TP): _____
- False Positives (FP): _____
- False Negatives (FN): _____
- True Negatives (TN): _____

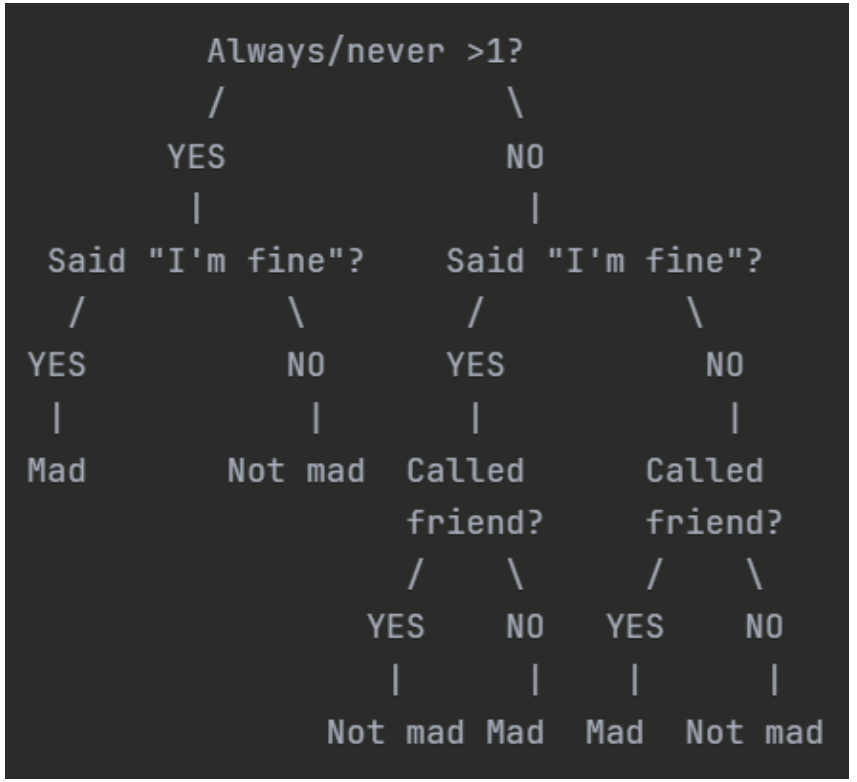


Confusion Matrix

	Predicted: Mad	Predicted: Not Mad
Actual: Mad	----- (True Positive)	----- (False Negative)
Actual: Not Mad	----- (False Positive)	----- (True Negative)

EVALUATE Tree 3

Row	Said "I'm fine"	"Always/ never" freq	Called friend	Actual : Mad	Model Predicts	Result
1	Yes	>1	Yes	Mad		
2	Yes	>1	No	Mad		
3	Yes	≤1	Yes	Not mad		
4	Yes	≤1	No	Mad		
5	No	>1	Yes	Not mad		
6	No	>1	No	Not mad		
7	No	≤1	Yes	Mad		
8	No	≤1	No	Not mad		



- True Positives (TP): _____
- False Positives (FP): _____
- False Negatives (FN): _____
- True Negatives (TN): _____

Confusion Matrix

	Predicted: Mad	Predicted: Not Mad
Actual: Mad	----- (True Positive)	----- (False Negative)
Actual: Not Mad	----- (False Positive)	----- (True Negative)

Workshop Recap
Tracing Our Journey
Through
AI and Machine Learning