

**This PowerPoint has been created to be best viewed in
Microsoft Office 365.**

For seamless viewing experience, use Microsoft Office 365 or use the
PDF version provided.

MACHINE LEARNING ESSENTIALS

YOU ALWAYS WANTED TO KNOW

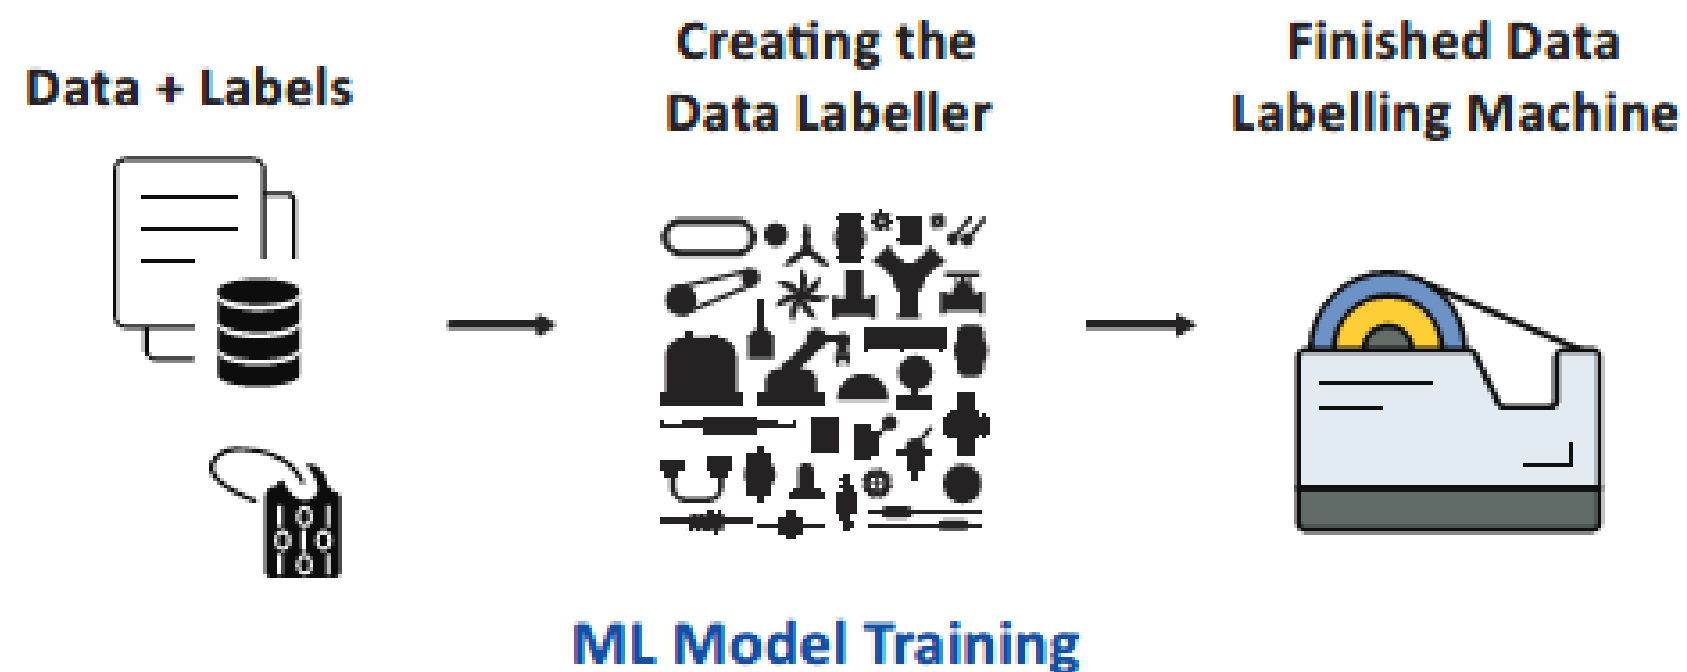


Author: Dhairya Parikh

CHAPTER 1: MACHINE LEARNING: A GENTLE INTRODUCTION

INTRODUCTION TO MACHINE LEARNING

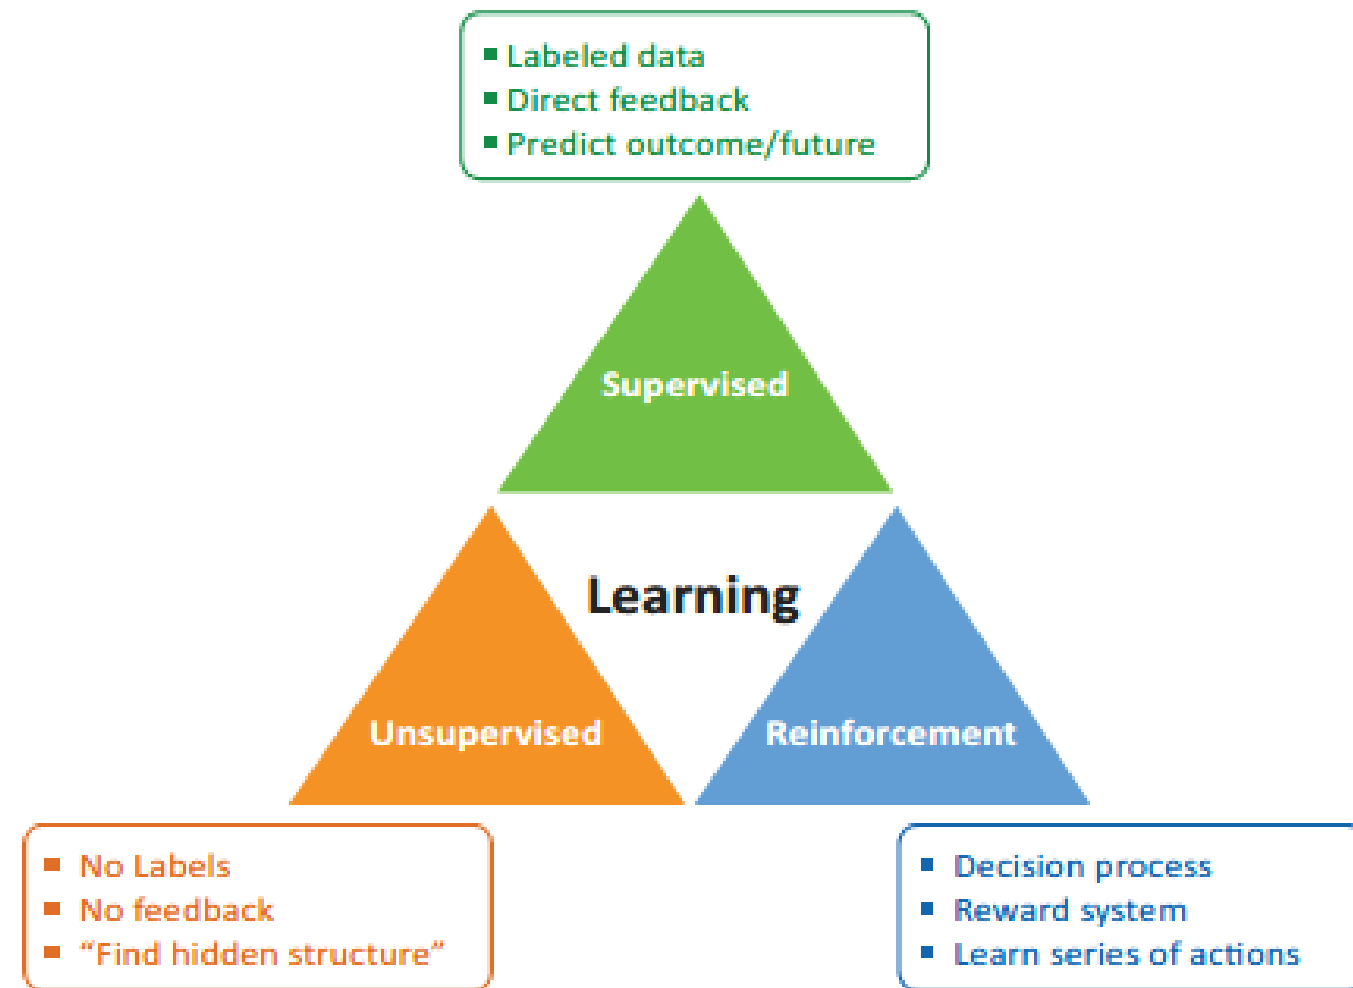
What is Machine Learning?



- A branch of AI focusing on creating machines that learn from data.
- Coined in 1959 by Arthur Samuel.
- Key Analogy: "Machine learning models are data labeling machines."

HOW MACHINES LEARN

Learning Paradigms in Machine Learning



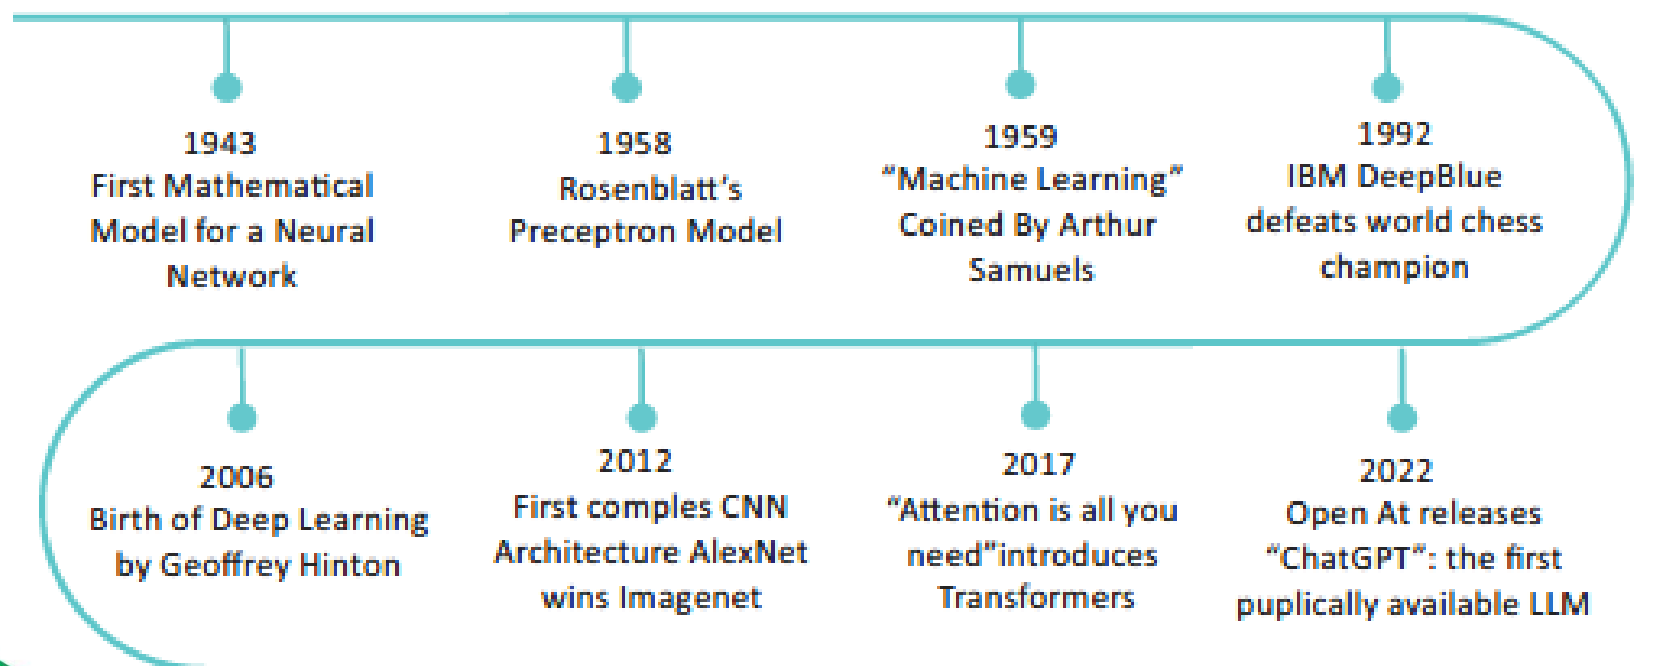
- Supervised Learning: Learns from labeled data.
- Unsupervised Learning: Finds patterns without labels.
- Reinforcement Learning: Learns by trial and error with rewards.

Source: Dhairya Parikh, "Learning Paradigms in Machine Learning." *DataDrivenInvestor*. Medium, June 15, 2020.
<https://medium.datadriveninvestor.com/learning-paradigms-in-machine-learning-146ebf8b5943>.

HISTORICAL EVOLUTION OF MACHINE LEARNING

Key Milestones in ML History


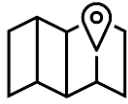

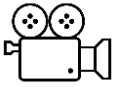
Machine Learning
An Historical Overview



- 1943: First Neural Network Model
- 1959: ML term coined.
- 1992: IBM's Deep Blue defeated a chess champion.
- 2012: Introduction of CNN with AlexNet.
- 2017: Google's Transformer model.

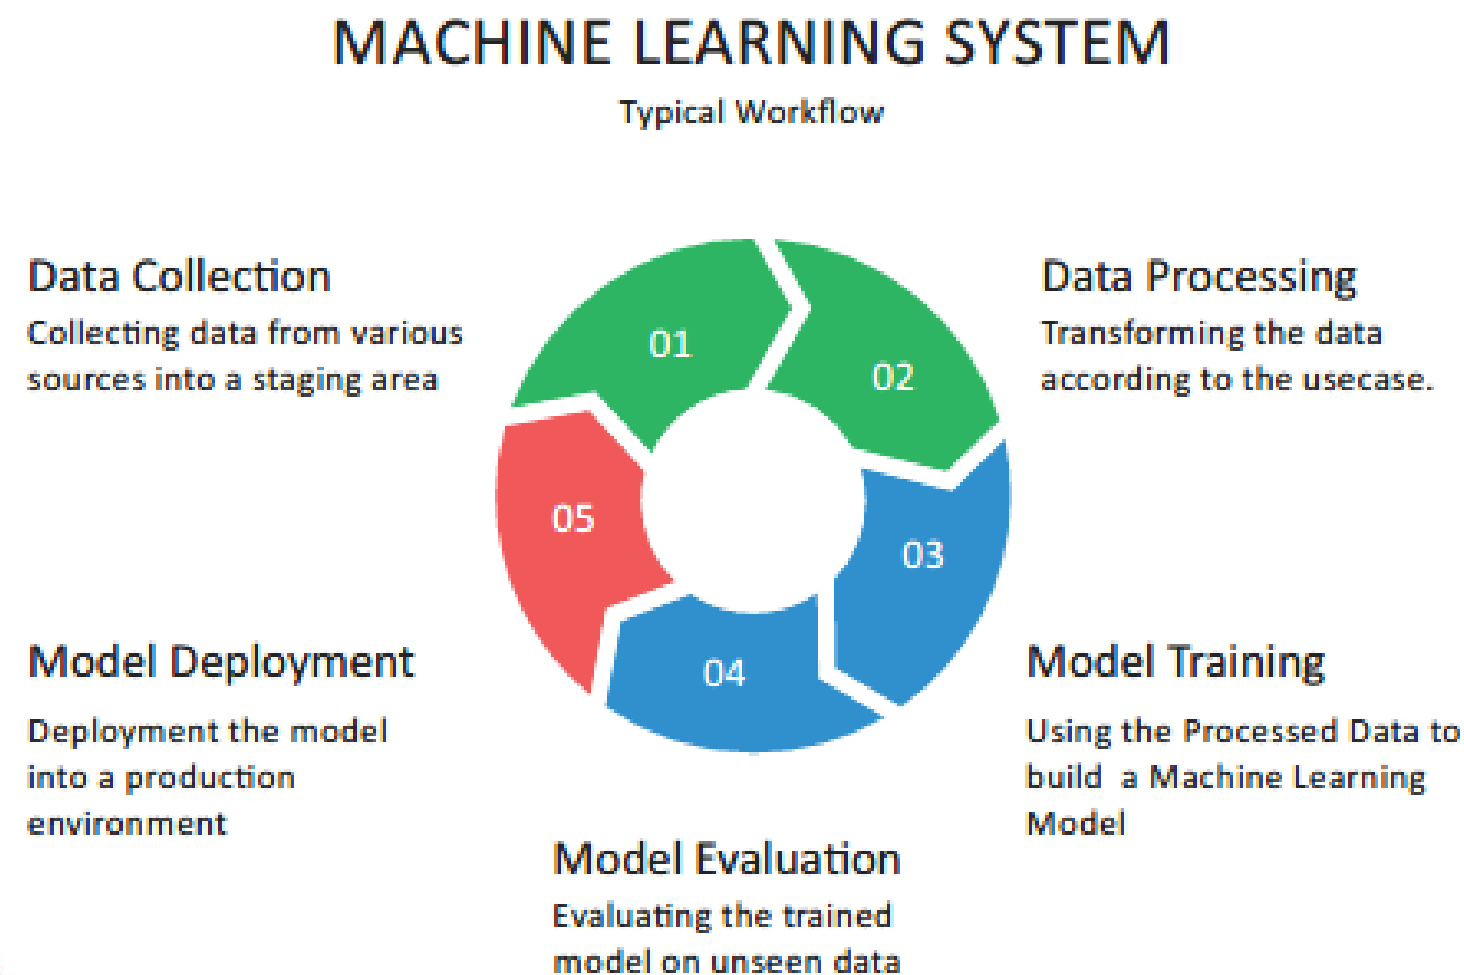
EVERYDAY APPLICATIONS OF MACHINE LEARNING

How ML Impacts Daily Life

- **Cell Phones:** Personalized content. 
- **Navigation:** Google Maps and traffic prediction. 
- **Personal Assistants:** Siri, Alexa, Google Assistant. 
- **Content Recommendation:** Netflix and Spotify's  algorithms.

COMPONENTS OF A TYPICAL ML SYSTEM

Building Blocks of ML Systems

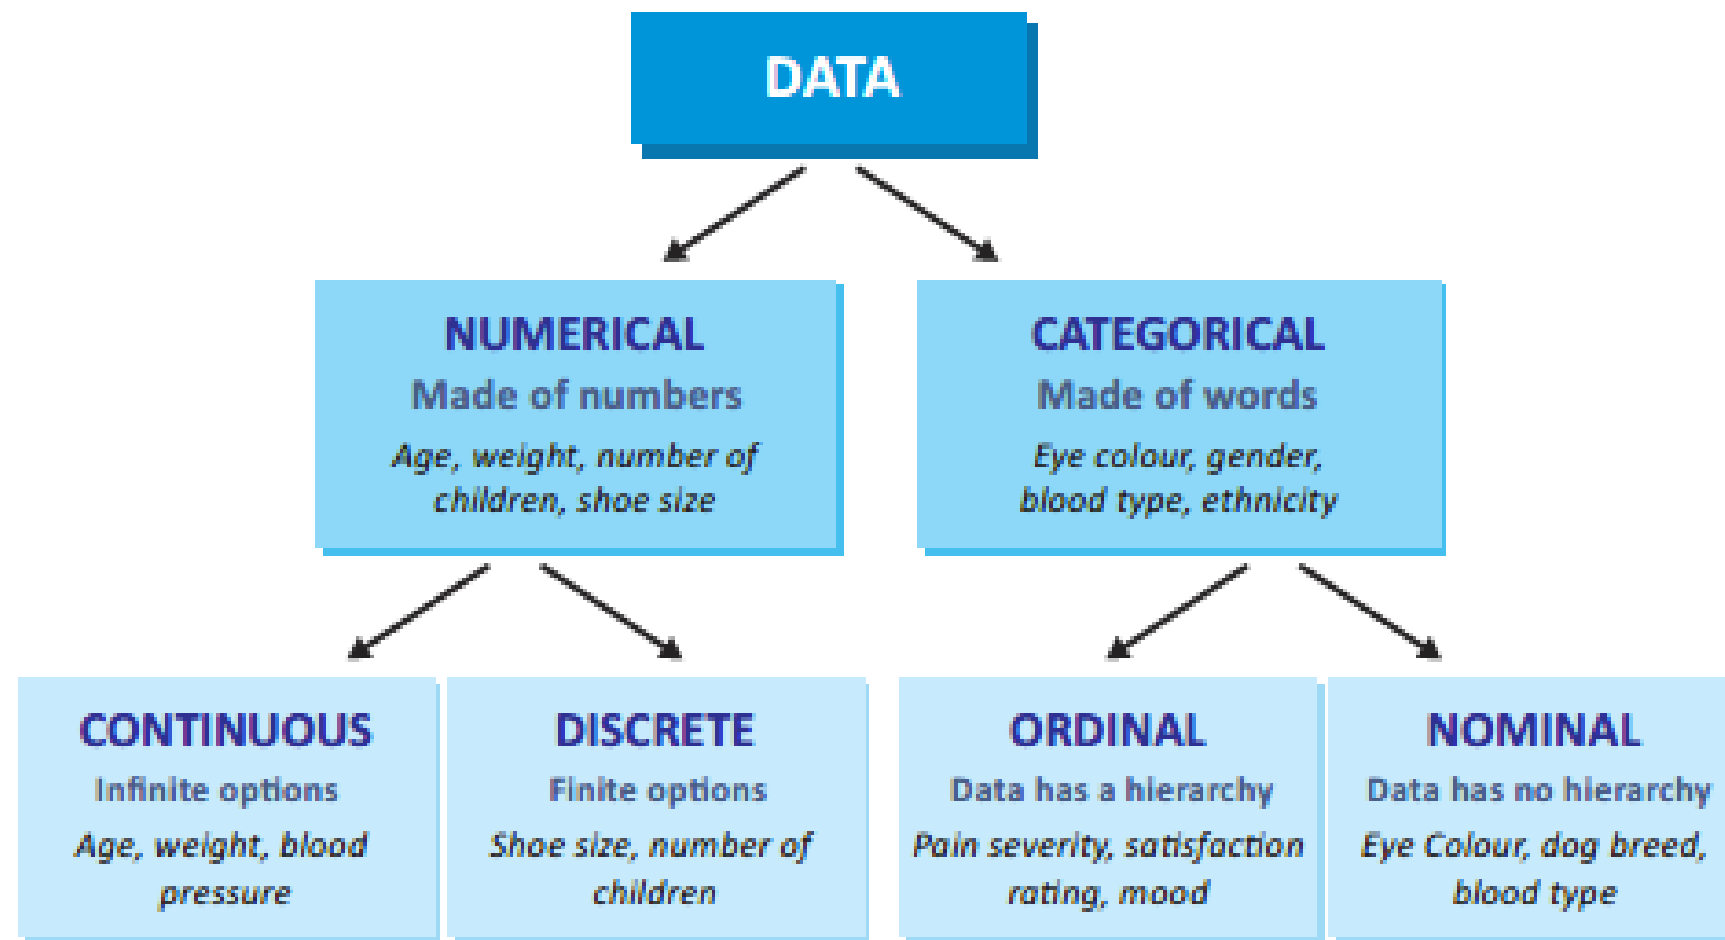


- **Data Collection:** Gathering data.
- **Data Processing:** Cleaning and organizing.
- **Model Training:** Learning from data.
- **Model Evaluation:** Testing accuracy.
- **Deployment:** Applying the model.

CHAPTER 2: MASTERING THE FUNDAMENTALS OF MACHINE LEARNING


- Importance of data, math, and programming in ML.
- **Overview:** Data types, essential math, Python, and key libraries.
- **Goal:** To equip you to handle ML data, understand math concepts, and start coding.

TYPES OF DATA IN MACHINE LEARNING



- **Numerical Data:** Continuous (e.g., temperature) & Discrete (e.g., book count).
- **Categorical Data:** Nominal (unordered) & Ordinal (ordered).
- Key concepts: **Cardinality** (number of samples) & **Dimensionality** (number of features).

THE MATH BEHIND MACHINE LEARNING



- **Linear Algebra:** Vectors, matrices, transformations (e.g., dot product).
- **Calculus:** Derivatives, gradients for optimization.
- **Probability & Statistics:** Predictions, patterns, and data analysis.

Why is Python the Ideal Language for ML?



- Simple syntax, vast libraries, cross-platform compatibility.
- Active community with ongoing support.
- Numerous frameworks and tools like **NumPy, Pandas, Matplotlib.**

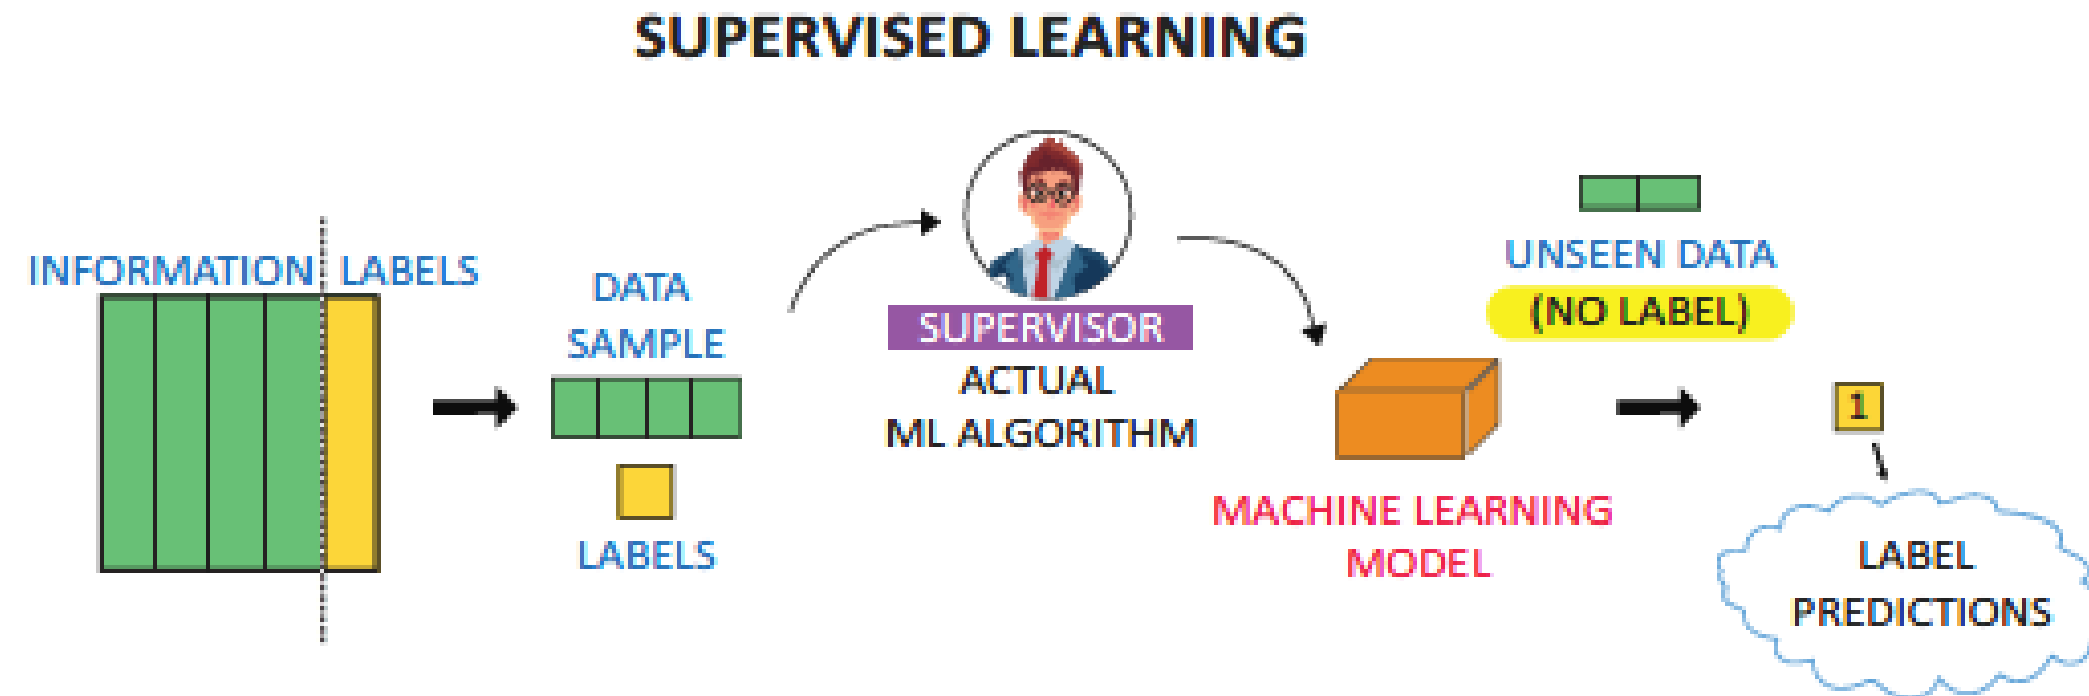
Introduction to Fundamental Python Libraries



- **NumPy:** Numerical computations with arrays and matrices.
- **Pandas:** Data manipulation using Data Frames.
- **Matplotlib:** Visualization from simple graphs to complex plots.

CHAPTER 3: SUPERVISED LEARNING: STARTING WITH THE BASICS


INTRODUCTION TO SUPERVISED LEARNING



- It involves training a model using labeled data.
- Key focus: Predicting outcomes from given data.
- We will explore **Linear Regression** for predicting continuous values and **Logistic Regression** for classification.

UNDERSTANDING LINEAR REGRESSION


Linear Regression in a Nutshell



- **Linear:** Predicts continuous values using a linear relationship.
- Key Concept: Find the best "line of fit."
- **Goal:** Minimize the error between predicted and actual values.

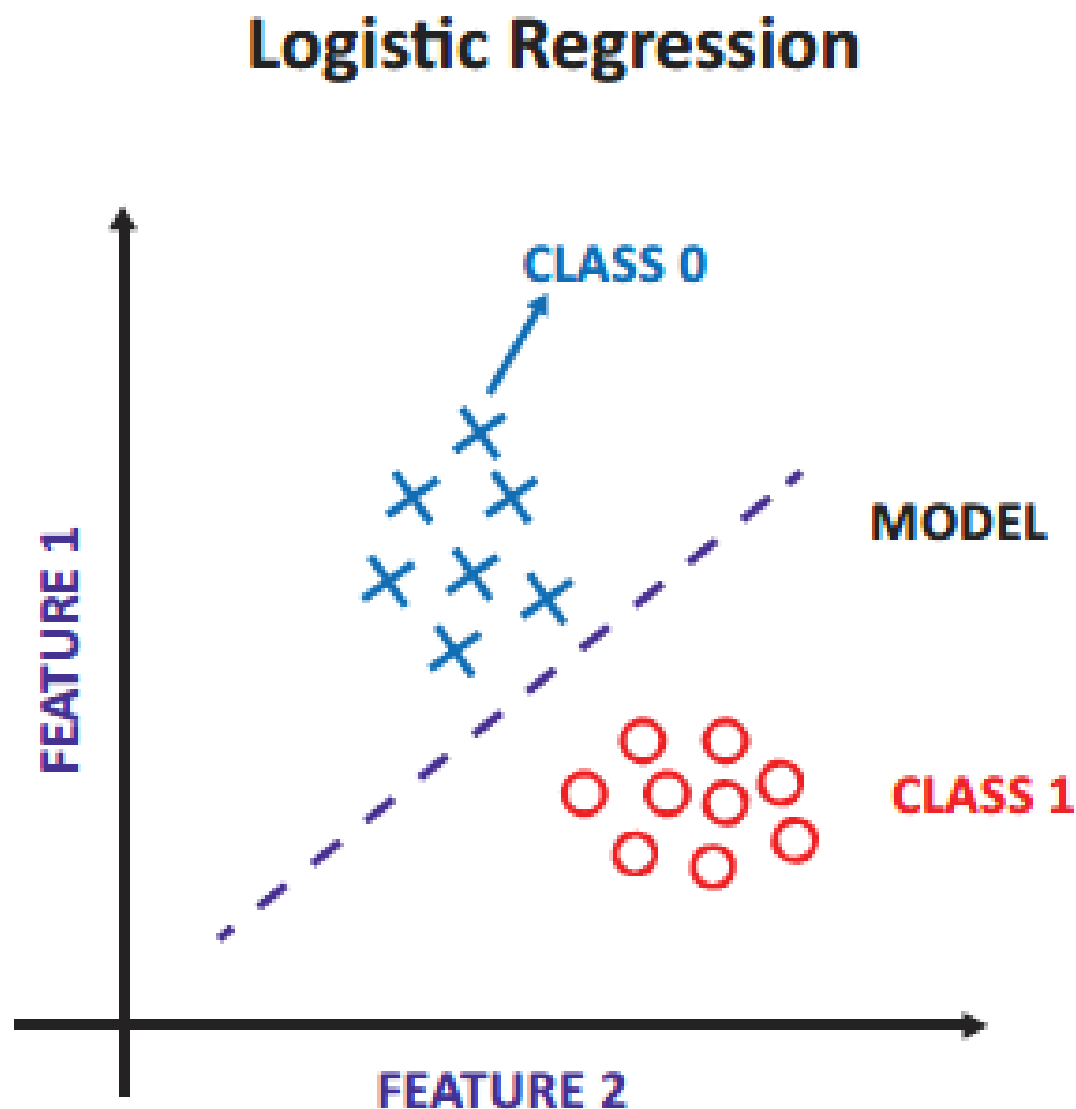
The Math Behind Linear Regression

Gradient Descent Algorithm



- **Model:** $y = w_0 + w_1x_1 + w_2x_2 + \dots + w_nx_n$
- Learning via minimizing error (Cost Function).
- Gradient Descent: Adjusting weights to reduce the error.

INTRODUCTION TO LOGISTIC REGRESSION



- Logistic Regression: For classification of tasks (binary/multi-class).
- Key Concept: Uses a **sigmoid function** to output probabilities.
- Goal: Find the decision boundary that best separates classes.

EVALUATING MODEL PERFORMANCE

$$\text{Accuracy} = \frac{\text{Number of Correct Predictions}}{\text{Total Number of Predictions}}$$

$$\text{Precision} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Positives}}$$

$$\text{Recall} = \frac{\text{True Positives}}{\text{True Positives} + \text{False Negatives}}$$

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$


- Classification: **Accuracy, Precision, Recall, F1-Score.**
- Regression: **MAE, MSE, RMSE, R-squared.**
- Confusion Matrix for classification insights.

CHAPTER 4:

GOING BEYOND THE BASICS: EXPLORING NON-LINEAR MODELS

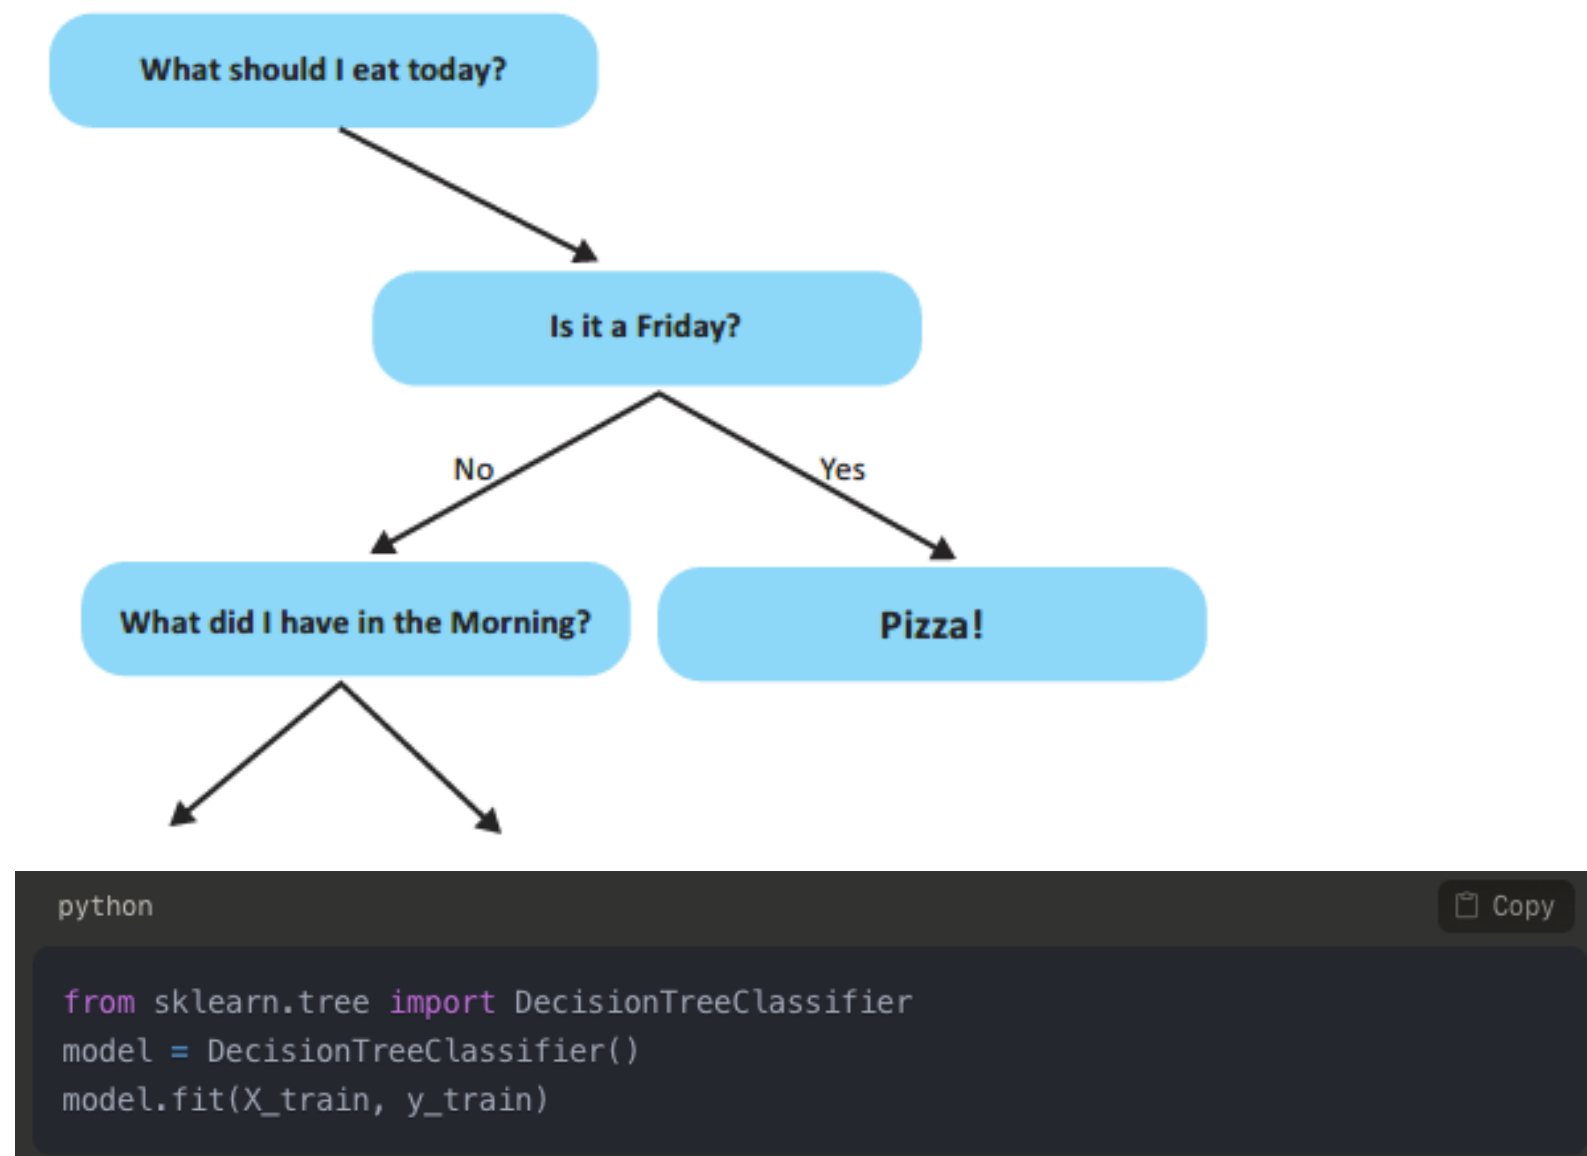
Introduction to Non-Linearity to the Equation

Linear vs. nonlinear problems



- Introduction to three powerful supervised learning algorithms.
- Focus on non-linear approaches to classification and regression
- Practical implementation in Python
- Hands-on evaluation and model selection.

DECISION TREES - MAKING HIERARCHICAL DECISIONS




- Tree-based approach uses if-else conditions
- Splits data based on feature values
- Creates a hierarchical structure for predictions
- Handles both classification and regression

K-NEAREST NEIGHBORS

Learning from Proximity

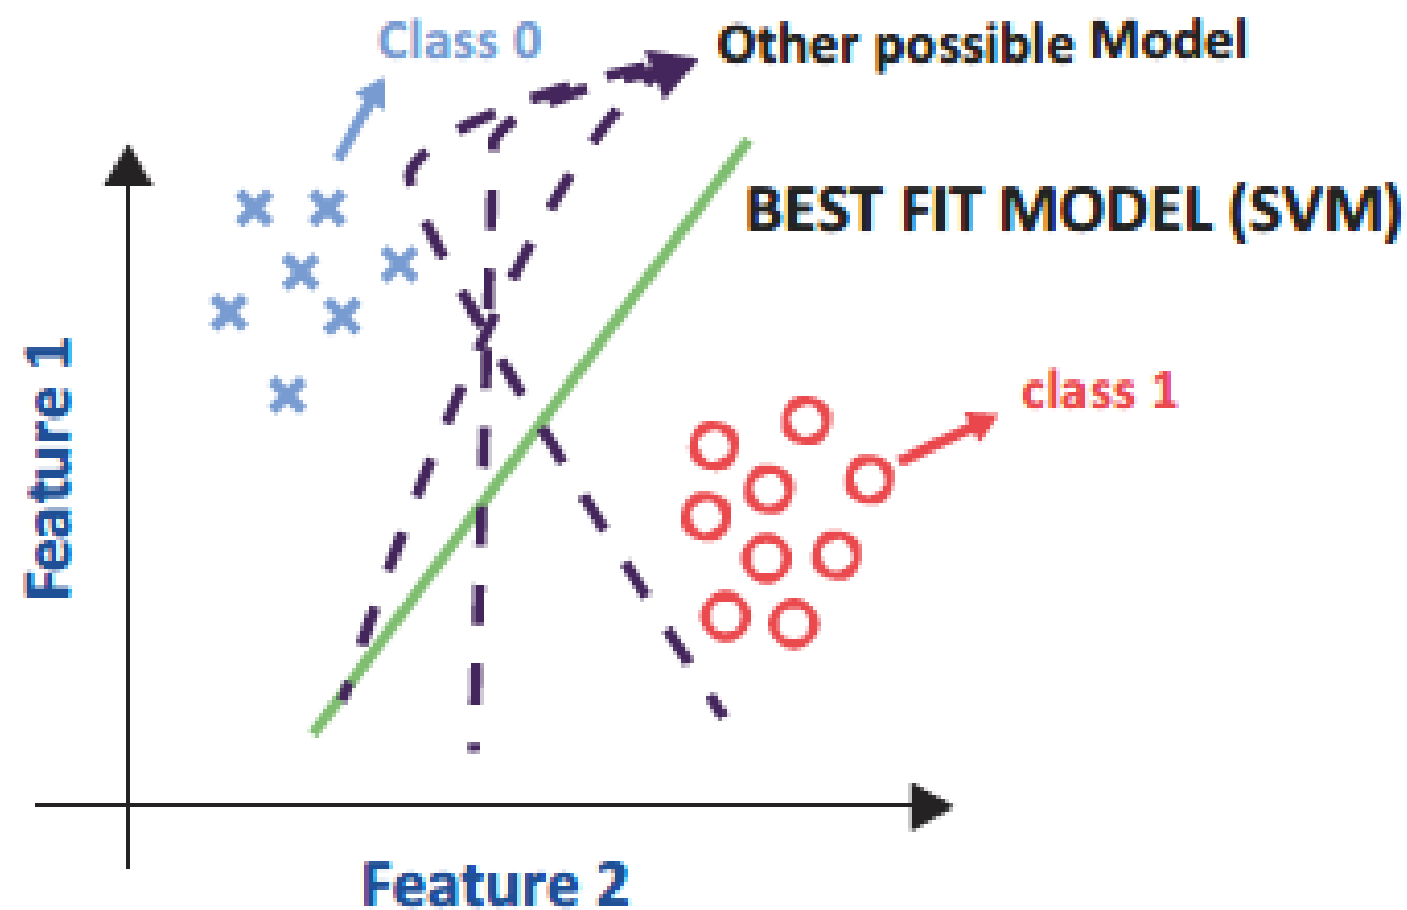
NEAREST NEIGHBOURS ALGORITHM



- Embrace simplicity by classifying data based on nearest examples
- Instance-based learning algorithm
- No explicit training phase
- Predictions based on closest neighbors
- Distance metrics determine similarity

SUPPORT VECTOR MACHINES

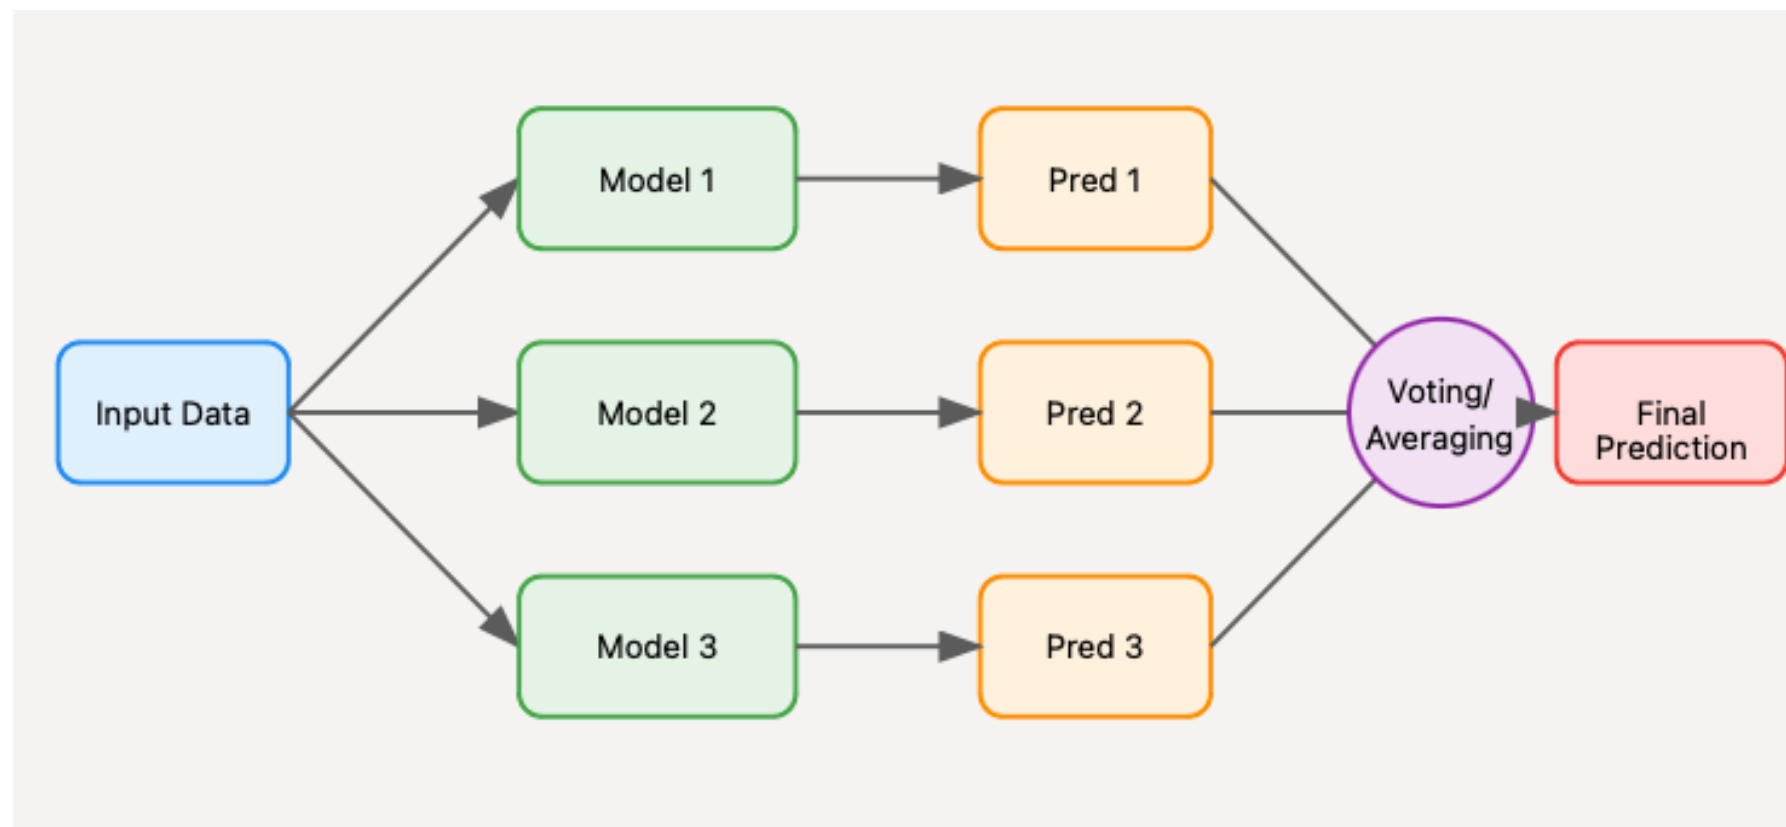
Finding Optimal Boundaries



- Master classification with margin maximization and kernel tricks.
- Maximum margin classification
- Kernel trick for non-linear problems
- Soft margin for handling outliers
- Support vectors define the boundary

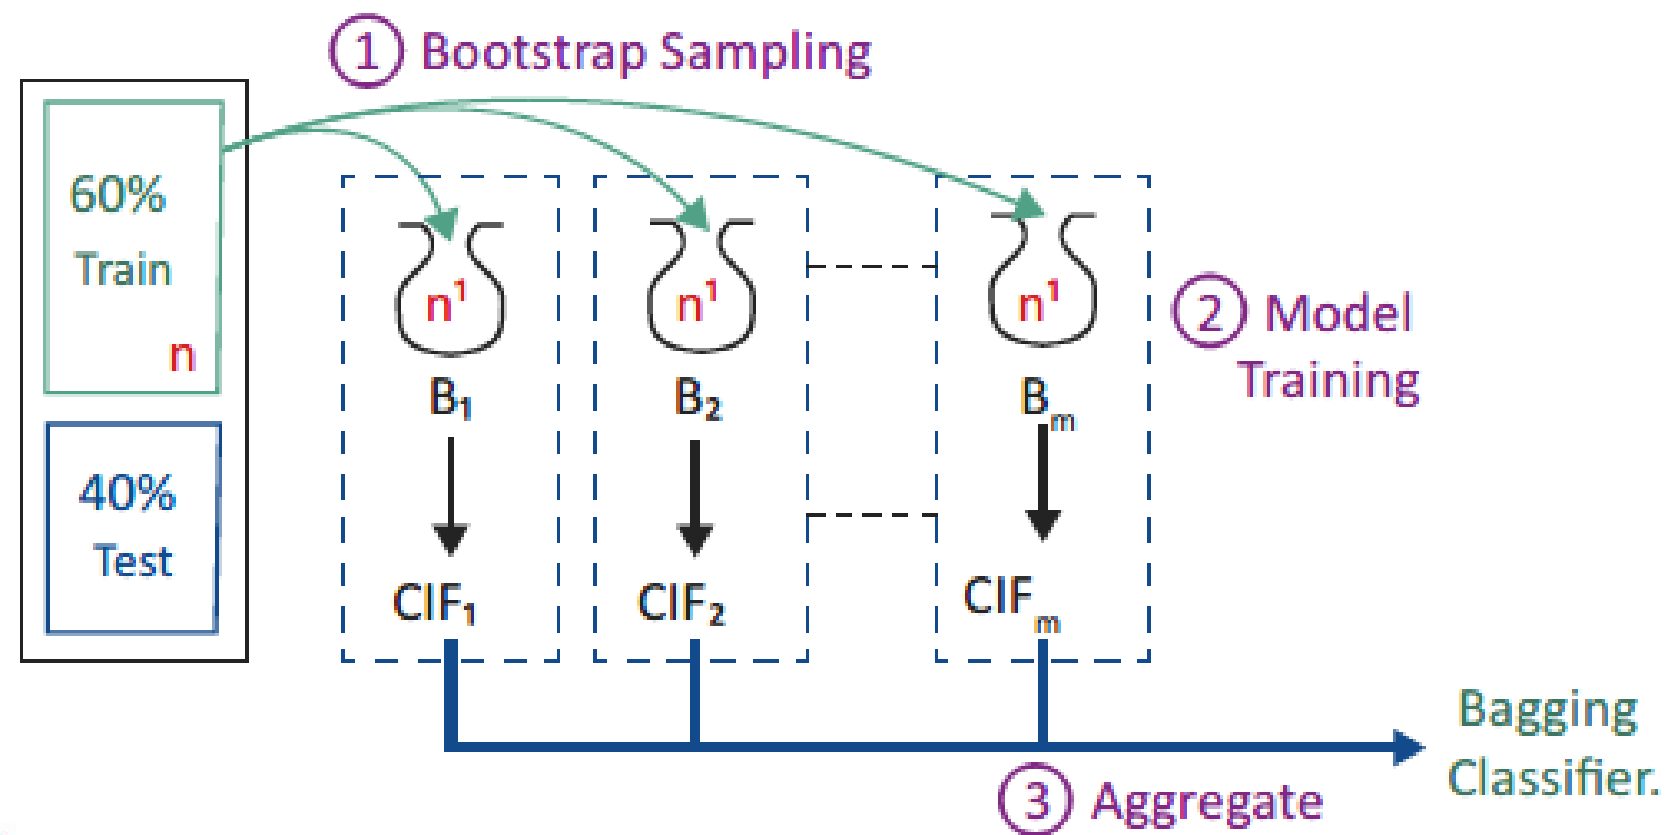
CHAPTER 5: ENSEMBLE TECHNIQUES: IMPROVING PREDICTION POWER

Why Use Ensemble Methods?



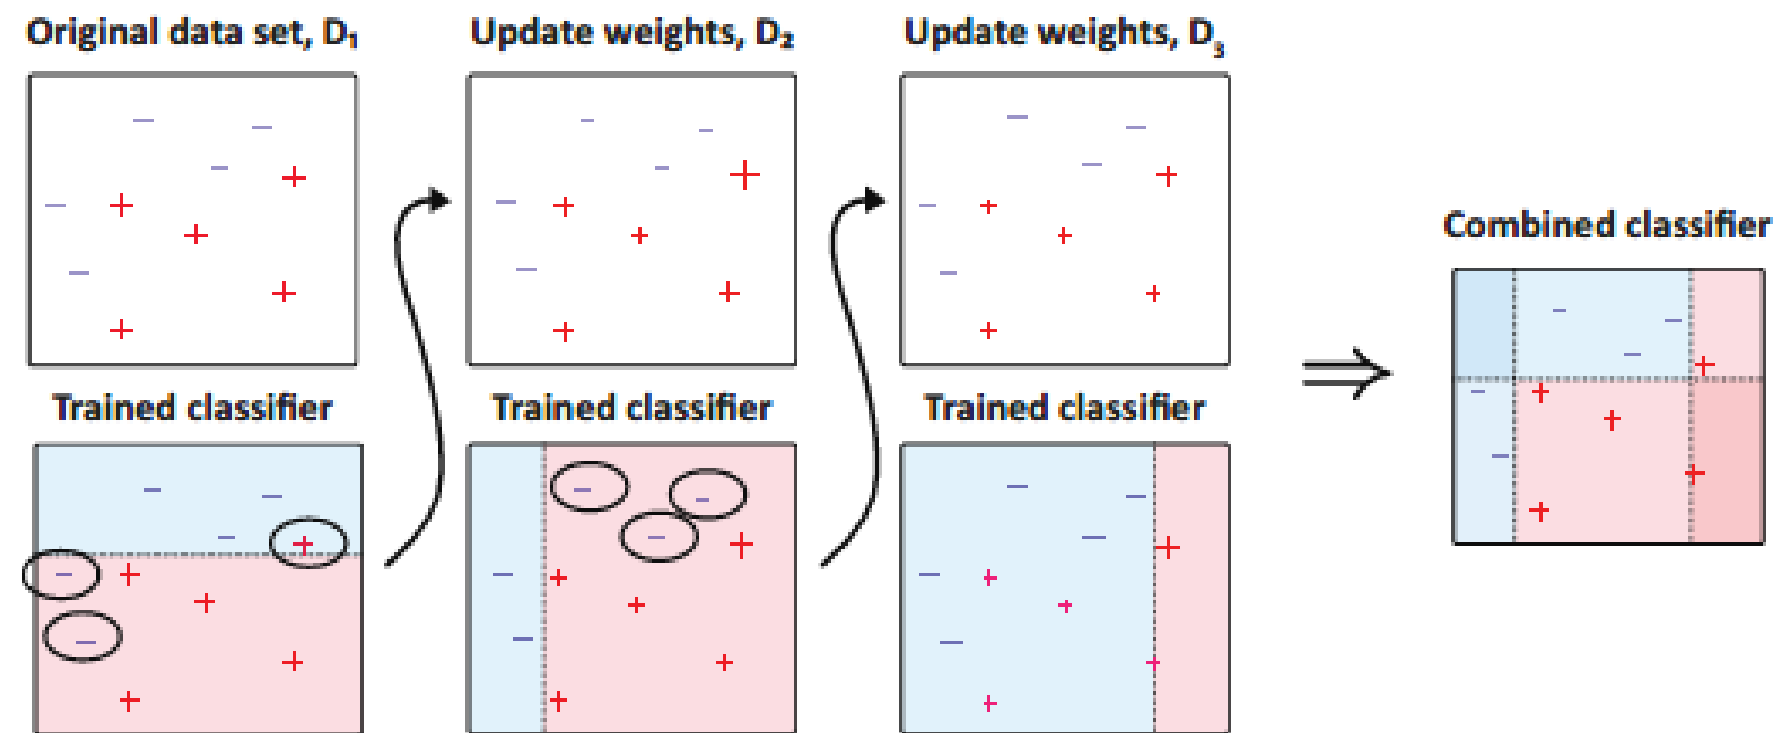
- Ensemble techniques combine multiple models to improve accuracy.
- Aim: Increase model robustness and reduce errors.
- Two main strategies:
Bagging and **Boosting**.

Bagging: Building Stronger Models



- **Definition:** Bootstrap Aggregating - multiple models trained on random data subsets.
- Key benefit: Reduces variance and prevents overfitting.
- Example algorithms: Random Forest.

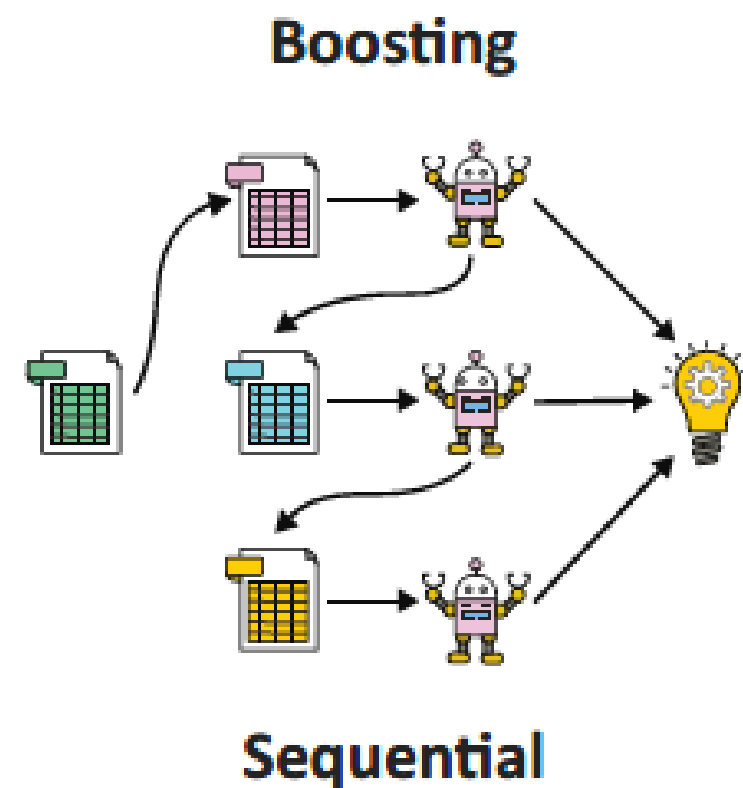

Boosting: Focus on Errors



- **Definition:** Sequential training of models where each corrects previous errors.
- Key benefit: Reduces bias and improves prediction accuracy.
- Example algorithms: AdaBoost, Gradient Boosting.

SUMMARY AND PERFORMANCE EVALUATION

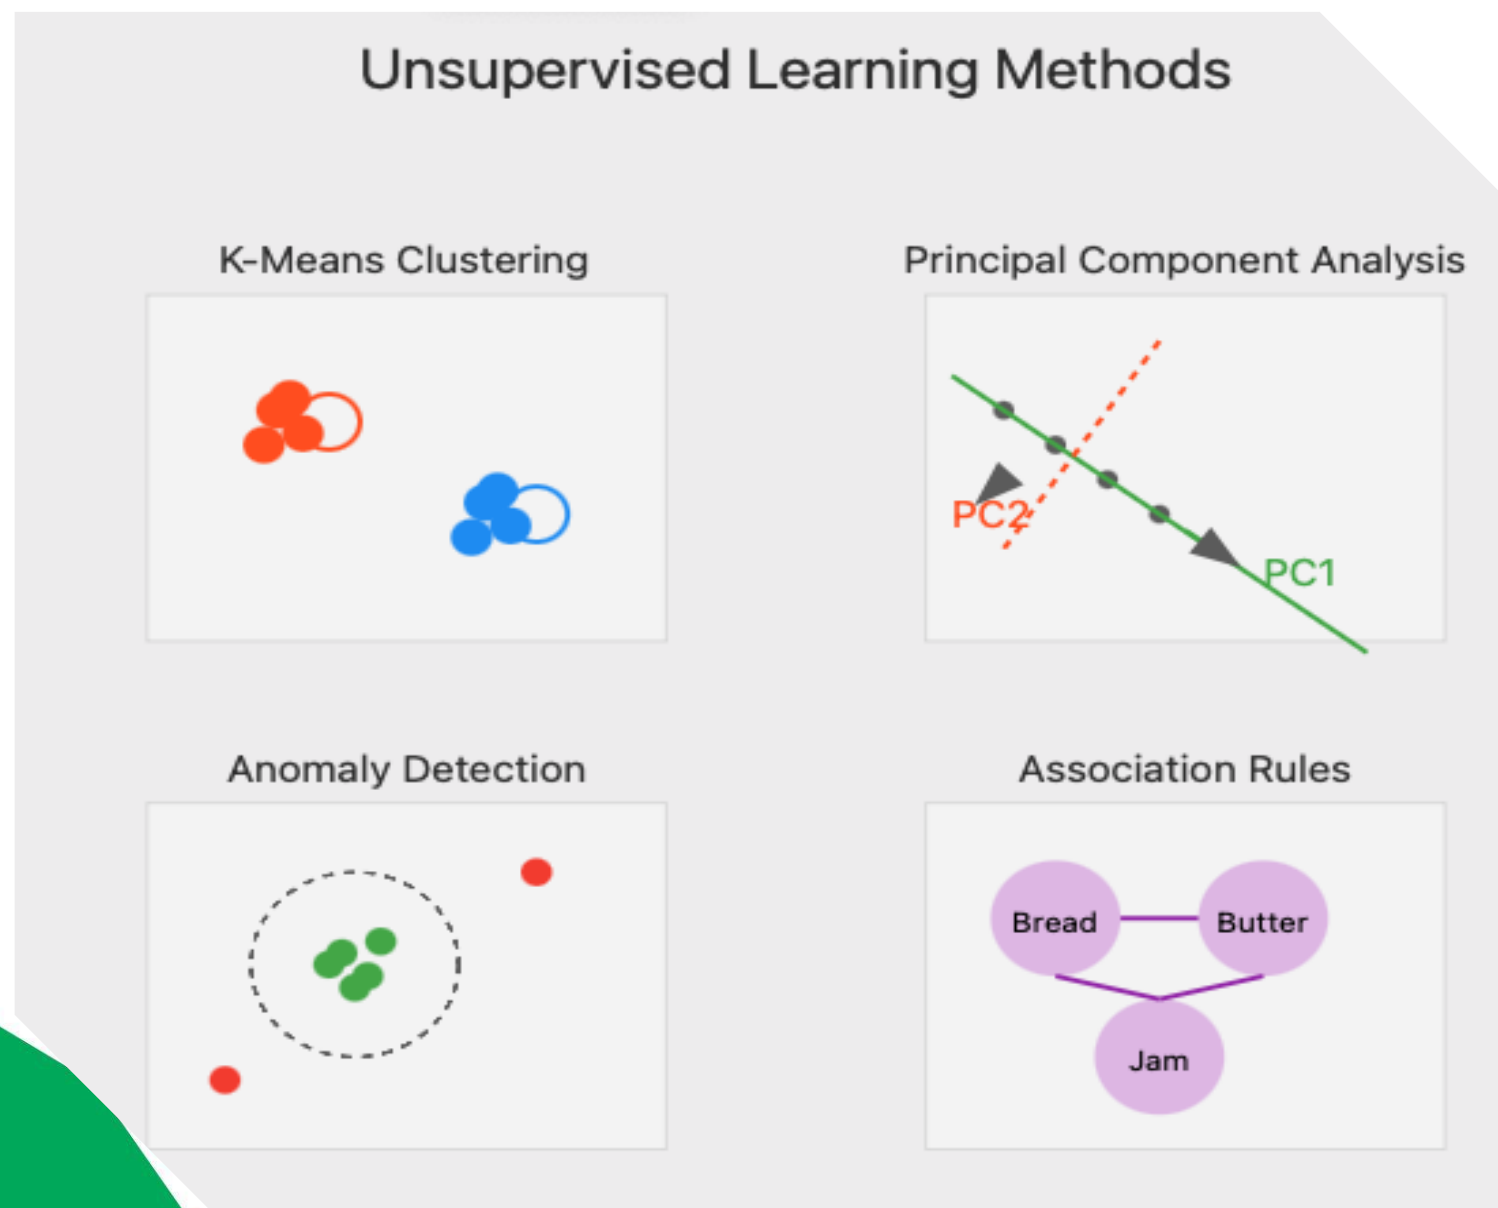
Choosing the Right Ensemble Method



- **Bagging:** Best for reducing variance and increasing stability.
- **Boosting:** Best for improving accuracy by reducing errors.
- **Evaluate models:** Compare using metrics like accuracy, precision, and recall.


CHAPTER 6: UNSUPERVISED LEARNING: FINDING PATTERNS IN DATA

What is Unsupervised Learning?



- Deals with unlabeled data.
- Finds hidden patterns and structures without predefined outcomes.
- Key Applications: **Clustering, Dimensionality Reduction, and Market Basket Analysis.**

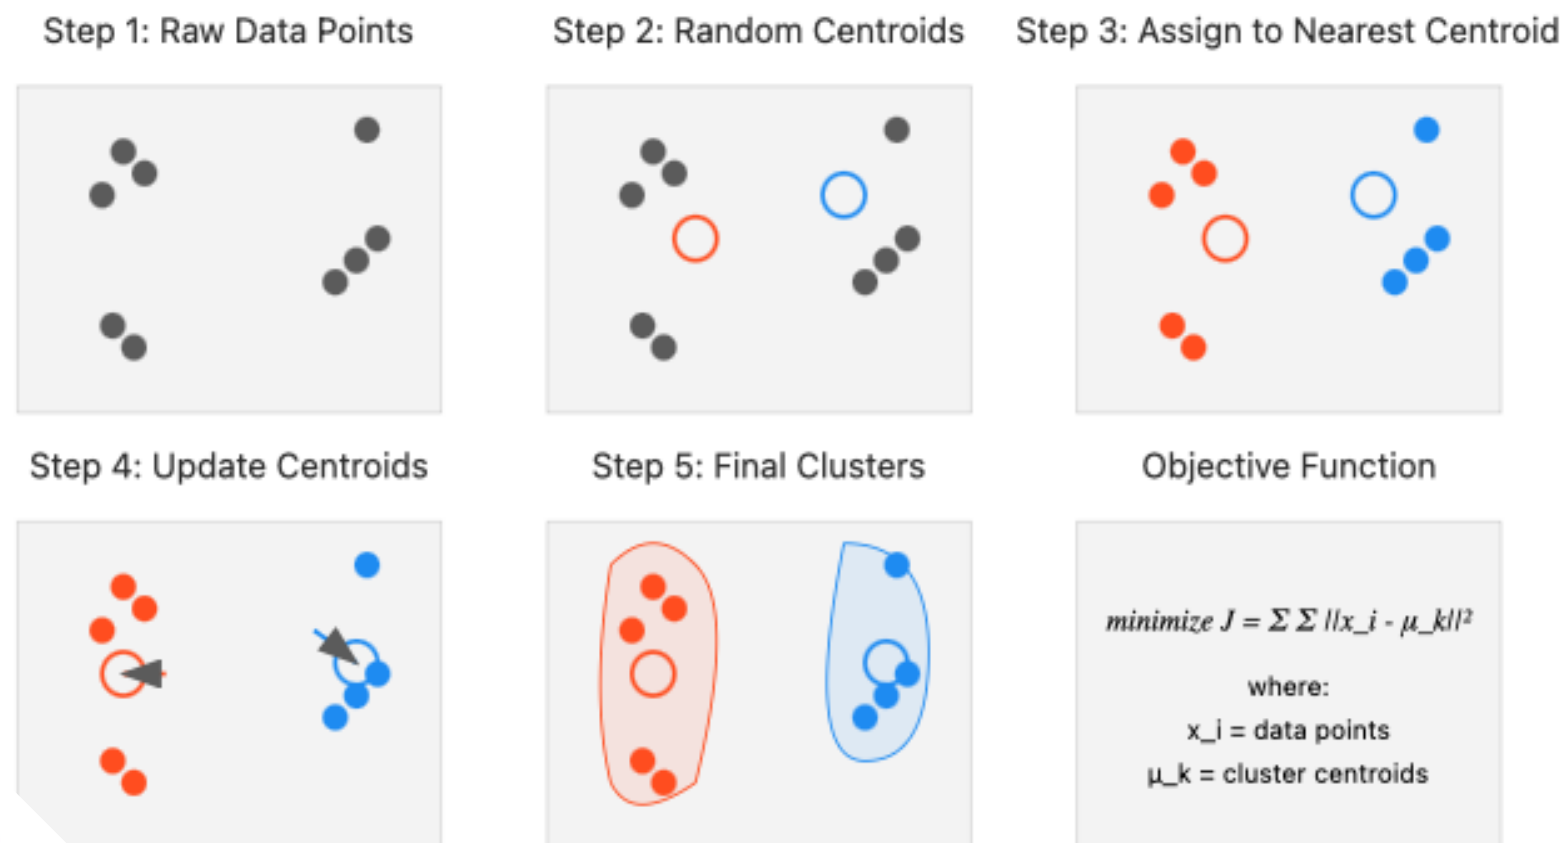
Clustering: Organizing Data into Groups



- Clustering groups of data points based on similarity.
- Useful for discovering patterns in large datasets.
- Key Algorithms: **K-Means, Hierarchical Clustering, DBSCAN.**

How K-Means Clustering Works

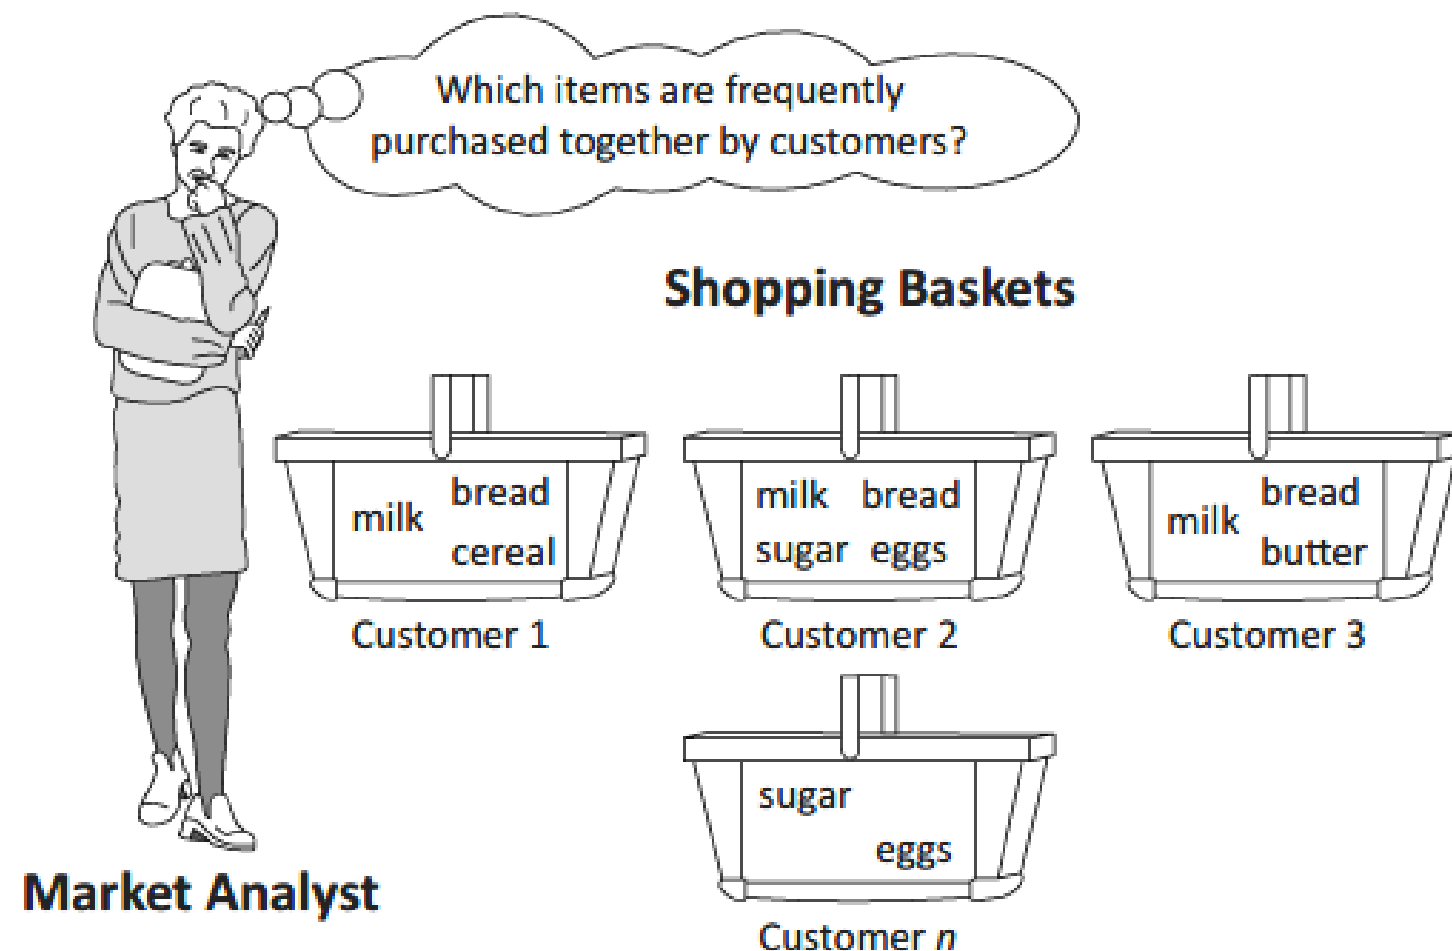
K-Means Clustering Algorithm Steps



- Divides data into K clusters based on proximity to centroids.
- Iteratively refines clusters to minimize intra-cluster distance.
- Pros: Simple, effective for large datasets.
- Cons: Sensitive to initial selection of centroids.

REAL-WORLD APPLICATION - MARKET BASKET ANALYSIS

Uncovering Patterns in Retail Data




- Identifies associations between products in transaction data.
- Helps in recommendation systems and inventory management.
- Example: "Customers who bought X are likely to buy Y."

CHAPTER 7:

A GENTLE INTRODUCTION TO NEURAL NETWORKS AND DEEP LEARNING


What is Deep Learning?



- Deep Learning is a subfield of ML inspired by the human brain.
- Uses neural networks to identify patterns in data.
- Popular in applications like image recognition, NLP, and AI models like ChatGPT.


A BRIEF HISTORY OF NEURAL NETWORKS

Biological Neuron versus Artificial Neural Network



- Inspired by **Biology**: Mimicking the human brain.
- **Key Milestones**:
 - 1943: McCulloch-Pitts Model of a Neuron.
 - 1980s: Backpropagation algorithm for training.
 - 2006: Emergence of Deep Learning with more complex architectures.

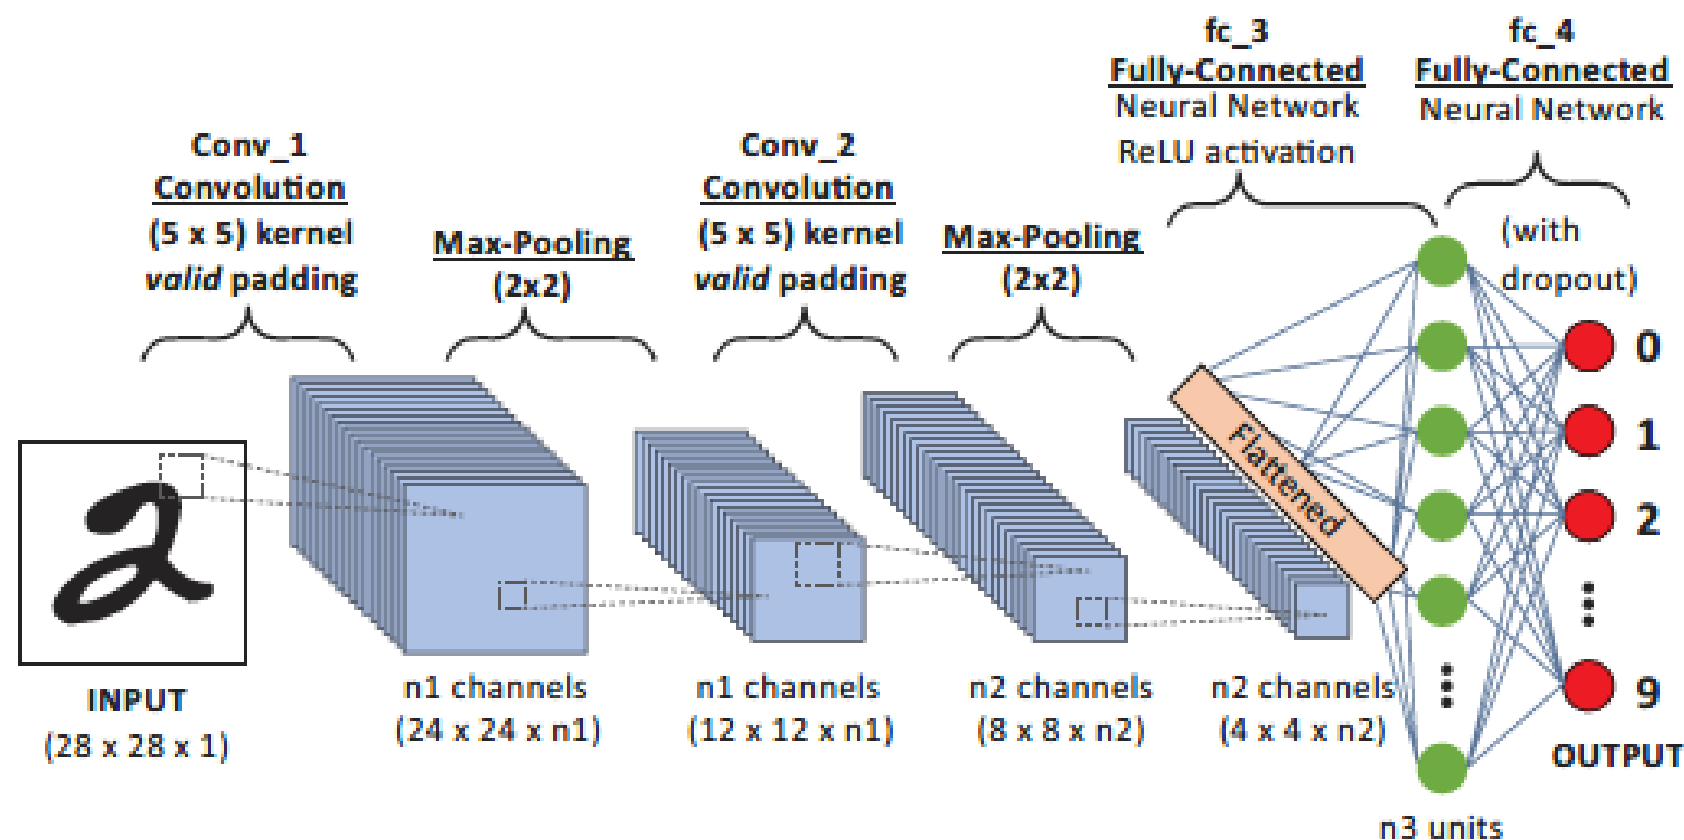
BUILDING DEEP NEURAL NETWORKS (DNNs)



- A network consists of layers: Input, Hidden, and Output.
- Deep Neural Networks have multiple hidden layers.
- Applications: Image recognition, speech processing, recommendation systems.

CONVOLUTIONAL NEURAL NETWORKS (CNNs)

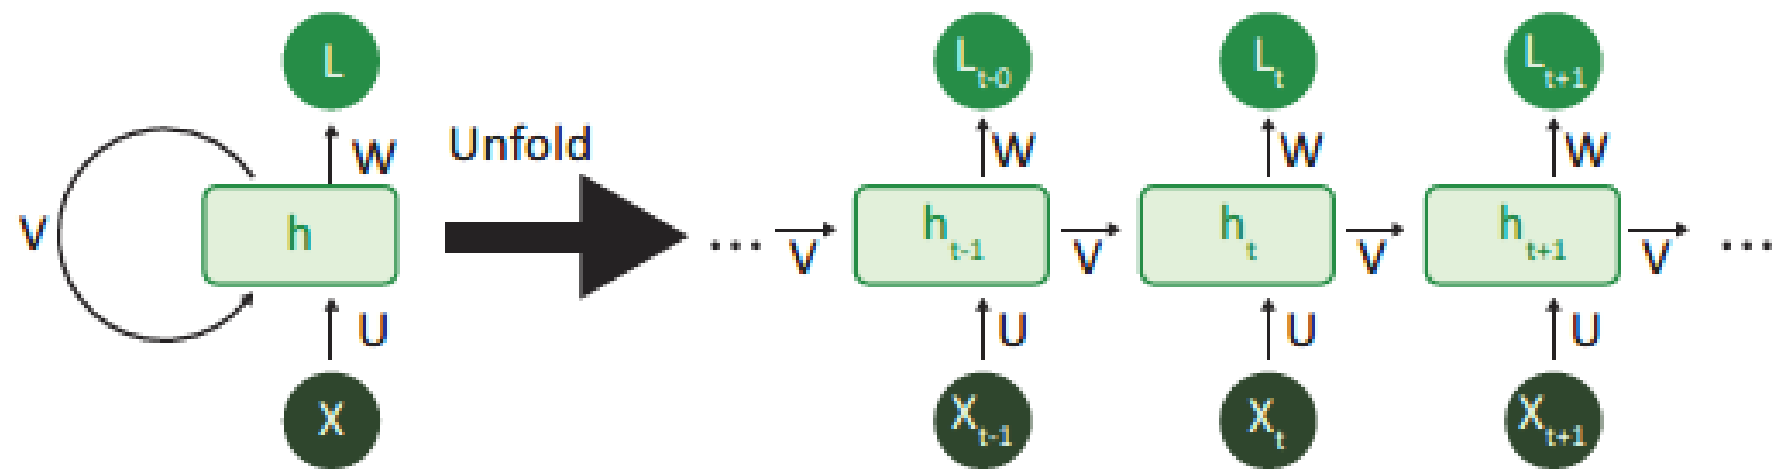
CNNs for Image Data



- Designed to handle **image and multimedia datasets**.
- Uses convolutional layers to extract features like edges, textures.
- Key advantage: Reduces computation while improving accuracy.

RECURRENT NEURAL NETWORKS (RNNs)

Understanding Sequences

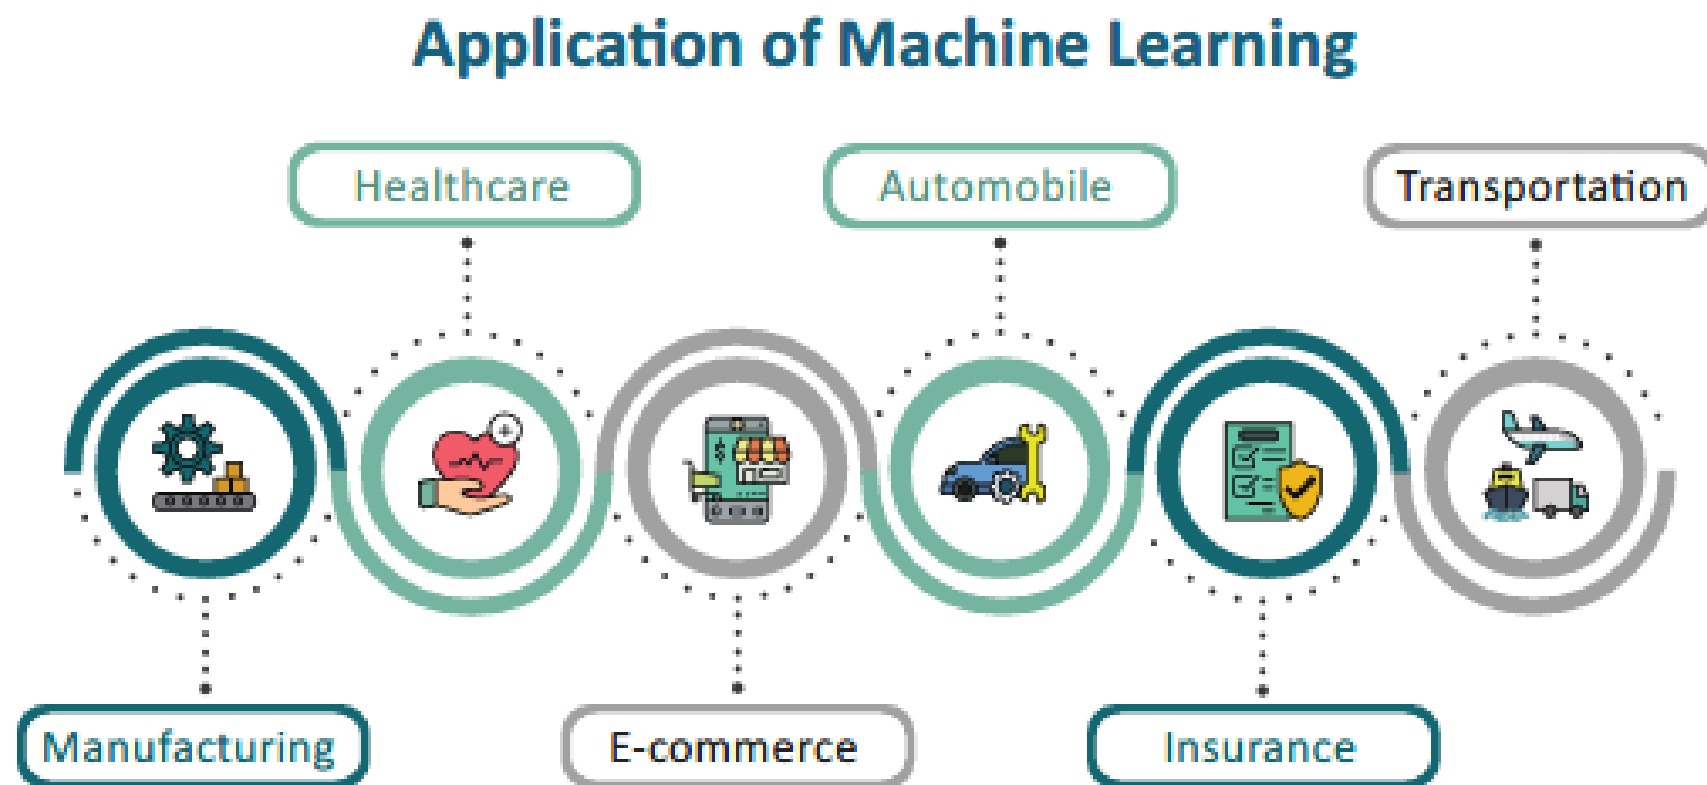


- Handles sequential data (e.g., text, time-series).
- Uses a feedback loop to consider previous information.
- Foundation for Large Language Models (LLMs) like ChatGPT.

CHAPTER 8: MACHINE LEARNING IN REAL-WORLD SCENARIOS

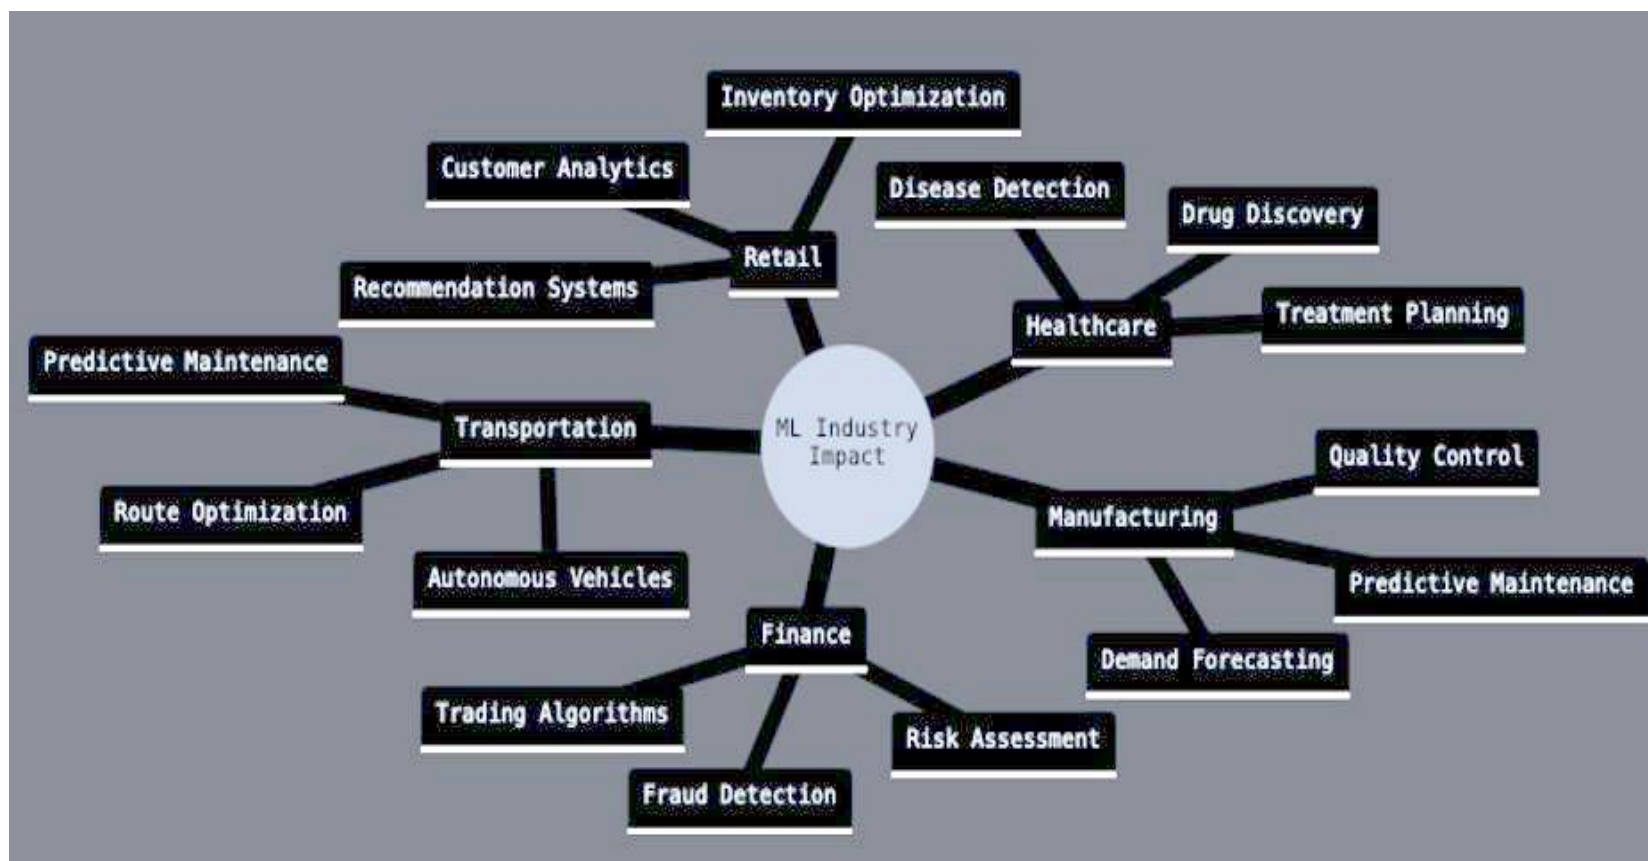
MACHINE LEARNING IN THE REAL WORLD

Applying Machine Learning Beyond the Classroom



- Recap of ML journey: From theory to real-world applications.
- Focus on industry use cases, project lifecycles, and ethical considerations.
- Goal: Bridge the gap between what you've learned and how to apply it.

How ML is Transforming Different Fields



- **Healthcare:** Early diagnosis, personalized treatment, predictive analytics.
- **Finance:** Fraud detection, credit scoring, algorithmic trading.
- **Hardware:** Predictive maintenance, optimization of manufacturing.
- Additional fields: Retail, marketing, autonomous vehicles.

THE ML PROJECT LIFECYCLE



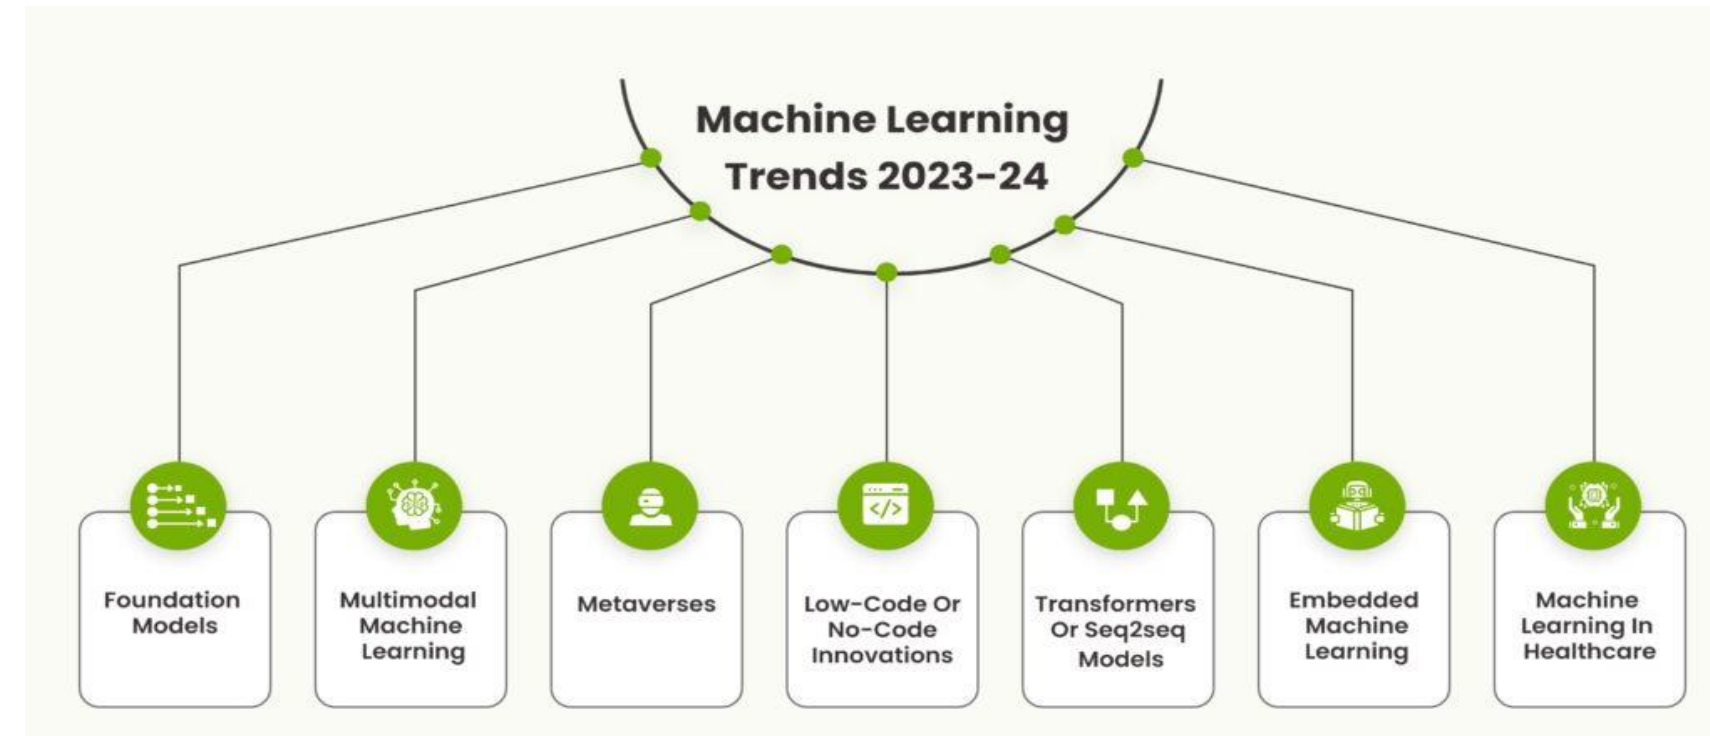
- **Step 1:** Problem Definition and Data Collection.
- **Step 2:** Data Cleaning, Processing, and Feature Engineering.
- **Step 3:** Model Training and Evaluation.
- **Step 4:** Deployment, Monitoring, and Maintenance.
- Importance of iterative improvement based on feedback.

ETHICAL CONSIDERATIONS IN MACHINE LEARNING



- **Fairness:** Ensuring unbiased models.
- **Transparency:** Making algorithms understandable.
- **Accountability:** Who is responsible for decisions made by AI?
- **Examples:** Bias in hiring algorithms, transparency in credit scoring.

FUTURE TRENDS IN MACHINE LEARNING



- **Emerging Technologies:** Federated Learning, Explainable AI, and Quantum Computing.
- **New Trends:** AI Ethics, Sustainable AI, TinyML (AI on low-power devices).
- **Organizations leading the way:** Open AI, Google DeepMind, MIT AI Lab.
- Emphasis on the growing importance of interdisciplinary collaboration.

Thank you