Introduction to Machine Learning

MODULE 1 Overview

Introduction

Ekpe Okorafor, PhD

Affiliations:

- **Accenture - Applied Artificial Intelligence**
  - ❑ Principal Director & Faculty, AI/ML & Knowledge Graph
- **African University of Science & Technology**
  - ❑ Visiting Professor, Computer Science / Data Science
  - ❑ Research Professor - High Performance Computing Center of Excellence

Research Interests:

- Big Data, Predictive & Adaptive Analytics
- Artificial Intelligence, Machine Learning
- Business & Information Systems
- Information Assurance and Cybersecurity.
- High Performance Computing & Network Architectures
- Distributed Storage & Processing
- Massively Parallel Processing & Programming
- Fault-tolerant Systems

Email: eokorafor@aust.edu.ng; ekpe.okorafor@gmail.com;
Twitter: @EkpeOkorafor; @Radicube

Modules for this topic

1. Overview: What is Machine learning
2. Categories of Machine Learning
3. Notations
4. Machine Learning Application Development Approach
5. Recommender Systems
6. Building a Recommender Engine

Learning Objectives

Upon successful completion of this topic, you will be able to:

- Define machine learning
- Describe the categories of machine learning
- Decide when to leverage Machine learning
- Discuss approaches to ML application development
- Differentiate between the ML approaches and motivations
- Build a simple recommender engine in R

Module 1

What is Machine Learning?

What is Machine Learning?

Machines are taking over!

What is Machine Learning?

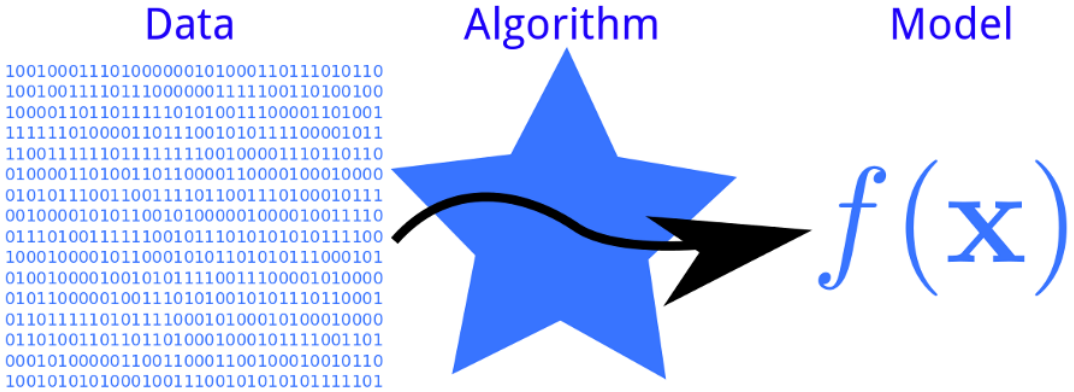

"Machine learning is the field of study that gives computers the ability to learn without being explicitly programmed. – Arthur L. Samuel, AI Pioneer, 1959

What is AI and ML?

- Artificial Intelligence (AI) is a branch or Computer Science that uses algorithms and techniques to mimic human intelligence
- Machine Learning (ML) is one of several AI techniques for sophisticated cognitive tasks

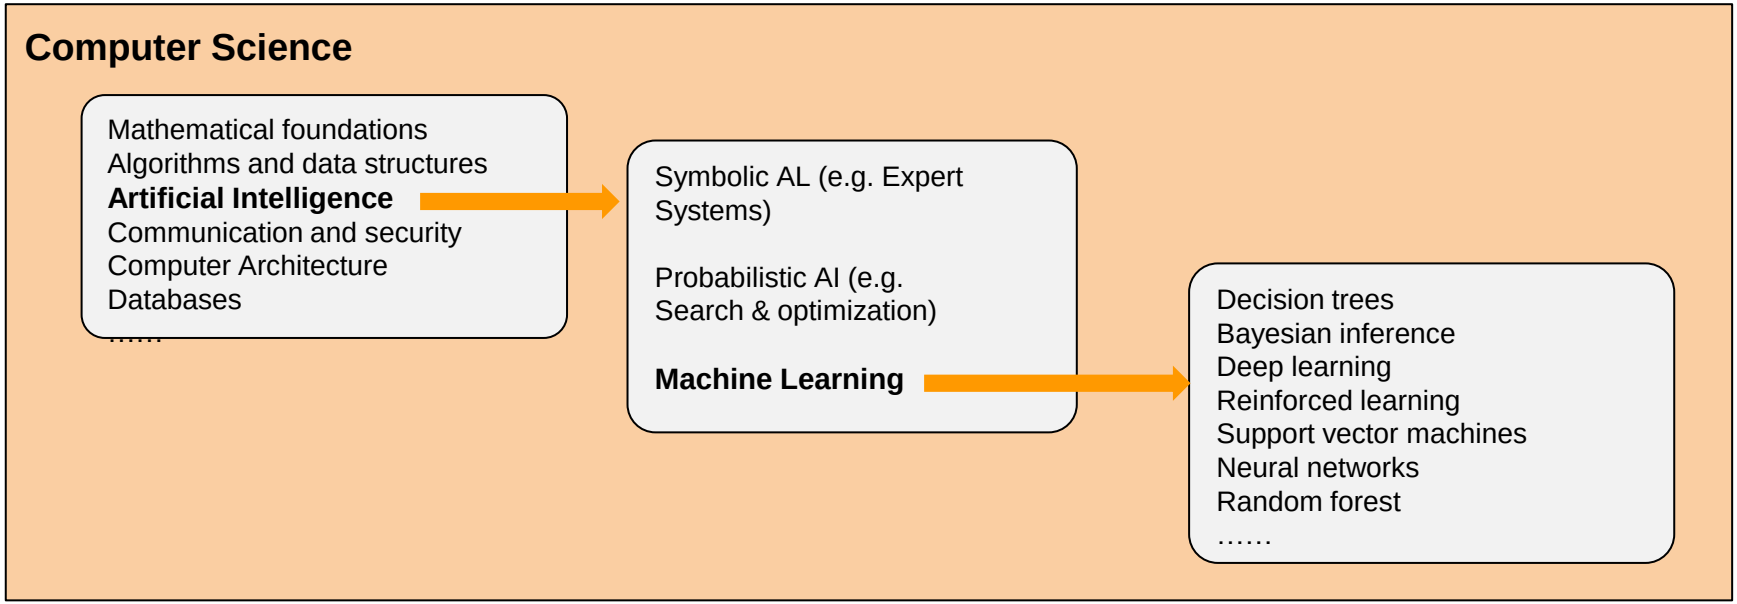

Machine Learning

Machine Learning is a particularly interesting technique because it represents a paradigm shift within AI

Traditional AI techniques

- **Static** – hard-coded set of steps and scenarios
- **Rule Based** – expert knowledge
- **No generalization** – handling special cases is difficult

---

Machine Learning

- **Dynamic** – evolves with data, finds new patterns
- **Data driven** – discovers knowledge
- **Generalization** – adapts to new situations and special cases

Machine Learning - Example

Example - Excelling at playing the game of chess

Symbolic AI	Mathematical/Statistical AI	Machine Learning Approach
“Let us sit down with the world’s best chess player, Ekpe Okorafor, and put his knowledge into a computer program”	“Let us simulate all the different possible moves and the associated outcomes at each single step and go with the most likely to win”	**“Let us show millions of examples or real life and simulated games (won and lost) to the program, and let it learn from experience”**

Machine Learning – When to use

- Machine learning is particularly good at solving 2 types of problems where other AI techniques fail

Tasks programmers can't describe

Complex multidimensional problems that can't be solved by numerical reasoning

Handwriting

Cognitive Reasoning

Weather Forecasting

Network Intrusion

Health Care Outcomes

Movie Recommendation

Applications of Machine Learning

- Email spam detection
- Face detection and matching (e.g., iPhone X, Windows laptops, etc.)
- Web search (e.g., DuckDuckGo, Bing, Baidu, Google)
- Sports predictions
- Post oce (e.g., sorting letters by zip codes)
- ATMs (e.g., reading checks)
- Credit card fraud
- Stock predictions
- Smart assistants (Apple Siri, Amazon Alexa, . . .)
- Product recommendations (e.g., Walmart, Netflix, Amazon)
- Self-driving cars (e.g., Uber, Tesla)
- Language translation (Google translate)
- Sentiment analysis
- Drug design
- Medical diagnoses
- ……

Exercise

As a first exercise, think about how machine learning could be applied in these problem areas or tasks listed in the previous slide:

- What is the desired outcome?
- What could the dataset look like?
- Is this a supervised or unsupervised problem, and what algorithms would you use? (Supervised and unsupervised learning will be introduced next)
- How would you measure success?
- What are potential challenges or pitfalls?

Summary

- We now have an overview and a good definition for Machine Learning
- We have seen some examples where ML can be leveraged