Machine Learning, Artificial Intelligence, and Natural Language Processing

Marcelo C. Medeiros

Department of Economics
The University of Illinois at Urbana-Champaign


WELCOME

Introduction

What is Machine Learning?

- A gentle introduction to
  1. Machine Learning (ML), Artificial Intelligence (AI), and
  2. Natural Language Processing (NLP)
- Open the black box of "Generative Pre-trained Transformers"
- Five online 90-minute meetings: 3/26, 4/2, 4/9, 4/16, and 4/23.

Introduction to the course

Do we care about these topics?

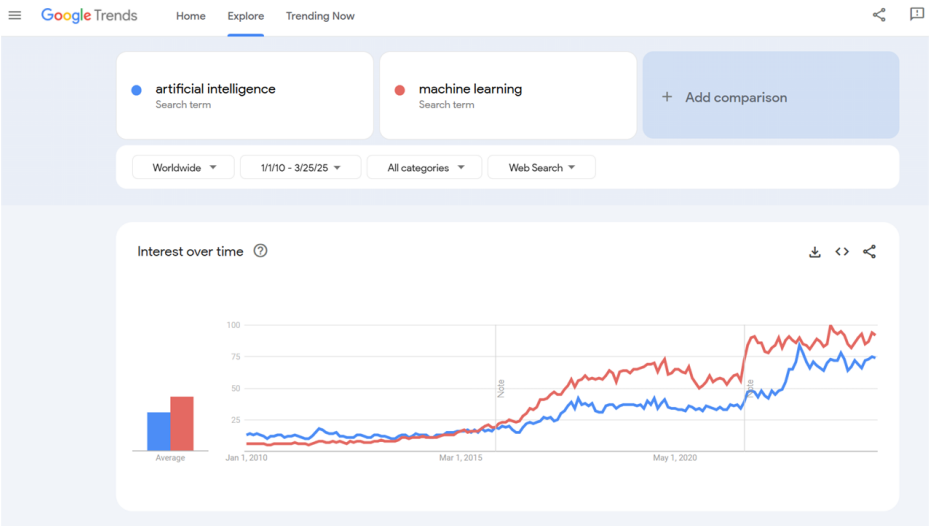

Introduction to the course

Do we care about these topics?

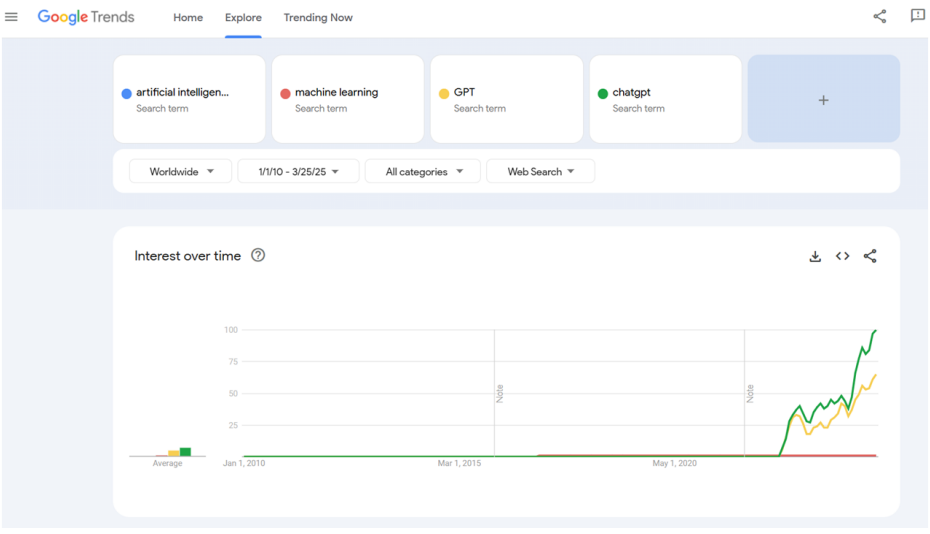

Introduction to the course

Do we care about these topics?

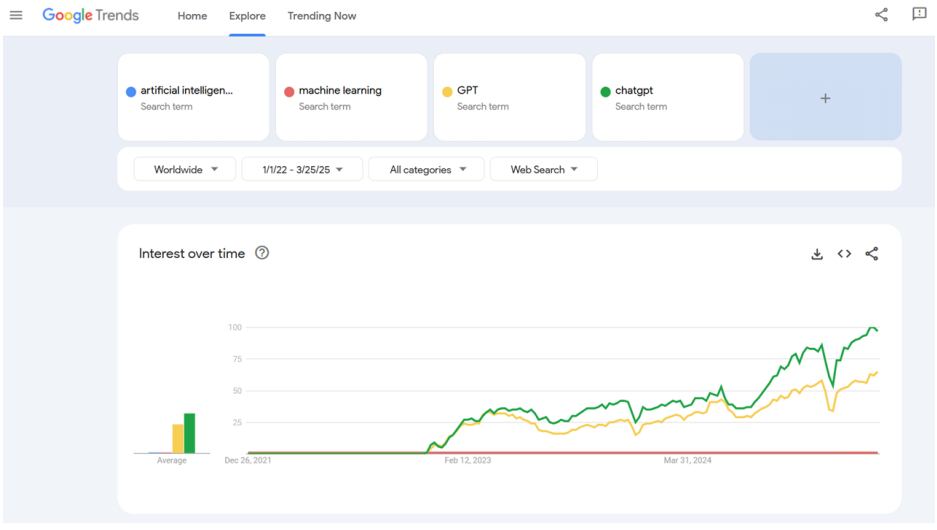

- ChatGPT is a chatbot based on generative pre-trained transformer (GPT) technology.

- ChatGPT is a chatbot based on generative pre-trained transformer (GPT) technology.
- GPTs are a family of large language models (LLMs) based on a transformer deep learning architecture.

- ChatGPT is a chatbot based on generative pre-trained transformer (GPT) technology.
- GPTs are a family of large language models (LLMs) based on a transformer deep learning architecture.
- It uses NLP and generative machine learning to craft dialogue that resembles human conversation.

Introduction to the course

What is ChatGPT?

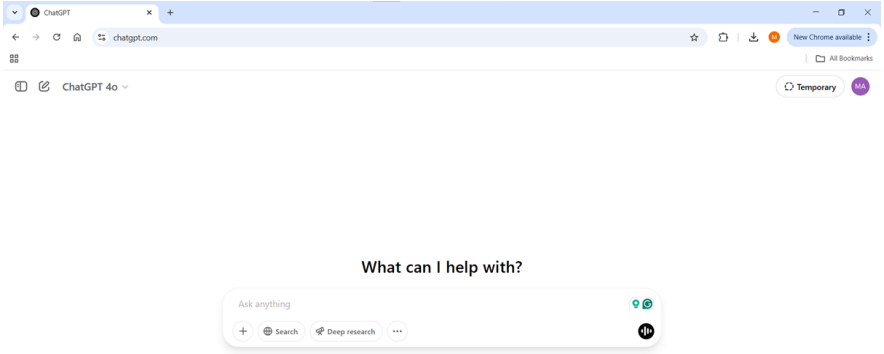

Introduction to the course

What is ChatGPT?

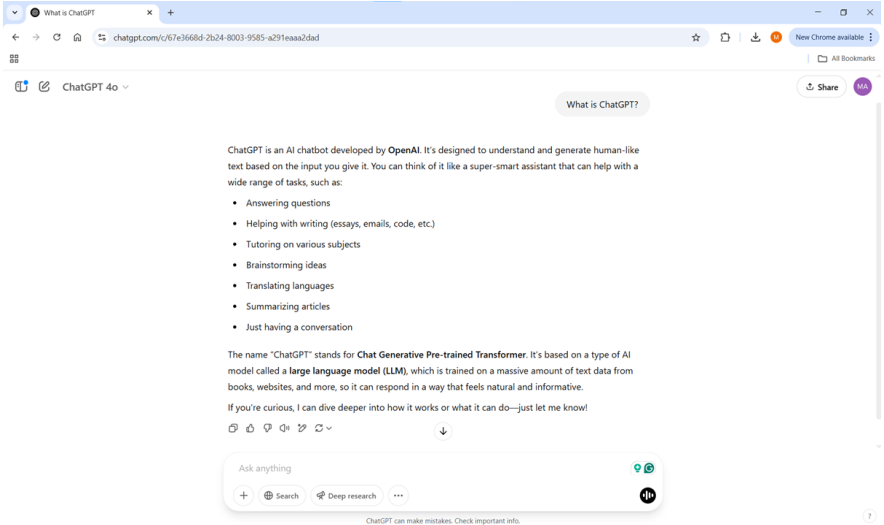

- There are many Large Language Models (LLM) out there ...
- Some examples:
  - Claude, BERT, RoBERTa, FinBERT, CoPilot, LLaMA, DeepSeek, ...
- To understand the "magic" behind the models we need to learn some basics of ML and NLP.

- What is a prompt?

- What is a prompt?
  - It is a natural language instruction that tells an AI model what to do.

- What is a prompt?
  - It is a natural language instruction that tells an AI model what to do.
  - It can be a query, command, or statement that includes context, instructions, and conversation history.

- What is a prompt?
  - It is a natural language instruction that tells an AI model what to do.
  - It can be a query, command, or statement that includes context, instructions, and conversation history.
- What is prompt engineering?

- What is a prompt?
  - It is a natural language instruction that tells an AI model what to do.
  - It can be a query, command, or statement that includes context, instructions, and conversation history.
- What is prompt engineering?
  - It is the practice of designing and refining prompts to get the desired output from an AI model.

- What is a prompt?
  - It is a natural language instruction that tells an AI model what to do.
  - It can be a query, command, or statement that includes context, instructions, and conversation history.
- What is prompt engineering?
  - It is the practice of designing and refining prompts to get the desired output from an AI model.
  - It is akin to offering a blueprint for the AI, directing it towards the precise result you envisage.

- Lecture 1: Introduction to ML and AI models
- Lecture 2: The basics of NLP
- Lecture 3: Neural Networks
- Lecture 4: Word embedding and attention mechanism
- Lecture 5: Neural Networks and LLMs – Transformers

What is Machine Learning?

Artificial Intelligence (AI) versus Machine Learning (ML)

- AI refers to the general ability of computers to emulate human thought and perform tasks in real-world environments.
- ML refers to the technologies (math) and algorithms that enable systems to identify patterns, make decisions, and *improve themselves through experience and data*.
- Machine learning is a pathway to artificial intelligence.

Natural Language Processing (NLP)

- NLP is a discipline that intersects computer science, linguistics, and artificial intelligence.
- NLP focuses on how computers comprehend, interpret, and handle human languages.
- It can involve things such as interpreting the semantic meaning of a language, translating between human languages, or recognizing patterns in human languages.
- It uses statistical methods, machine learning, and text mining.

Machine Learning, Artificial Intelligence, and Natural Language Processing

Lots of buzzwords...

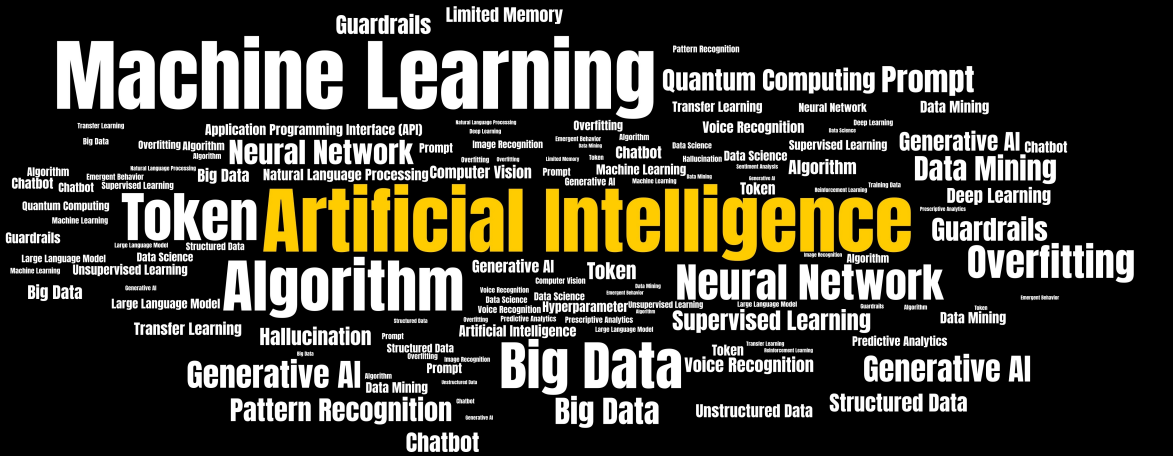

Machine Learning and Artificial Intelligence

AI and ML significantly influence our everyday activities...

navigation systems

social media

smart homes

mobile phones

assistants

chatbots

self-driving cars

recommendation systems

safety

language translation

- AI/ML methodologies have existed for quite some time...

- AI/ML methodologies have existed for quite some time...
- So, why have they gained such attention only recently?

- AI/ML methodologies have existed for quite some time...
- So, why have they gained such attention only recently?

improved fuel: huge amounts of data

- AI/ML methodologies have existed for quite some time...
- So, why have they gained such attention only recently?

improved fuel: huge amounts of data

advances in computational power and new methods/algorithms

- AI/ML methodologies have existed for quite some time...
- So, why have they gained such attention only recently?


improved fuel: huge amounts of data


advances in computational power and new methods/algorithms


the ability of process text data (and audio, images, ...)

We can find data everywhere…

Data come in very different formats…

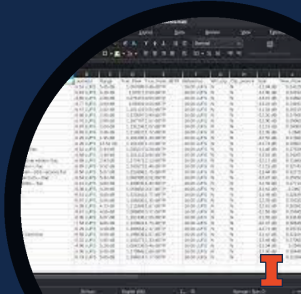

ECONOMETRICS LAB

Data arrive in large volumes….

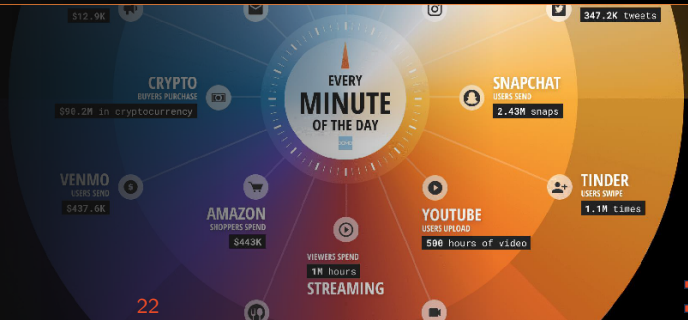

From Data to Decisions

Big Data

Machine Learning/AI

Decisions

What is Machine Learning?

What is Machine Learning?

Definitions out there...

Machine learning is the field of study that gives computers the ability to learn without being explicitly programmed.

Arthur L. Samuel, AI pioneer, 1959

The original sentence:

Programming computers to learn from experience should eventually eliminate the need for much of this detailed programming effort.

A.L. Samuel (1959). *Some studies in machine learning using the game of checkers.* **IBM Journal of Research and Development**, 3.3, 210–229.

What is Machine Learning?

Definitions out there...

- Machine Learning is the science of getting computers to learn and act like humans do, and improve their learning over time in autonomous fashion, by feeding them data and information in the form of observations and real-world interactions.

https://emerj.com/ai-glossary-terms/what-is-machine-learning/

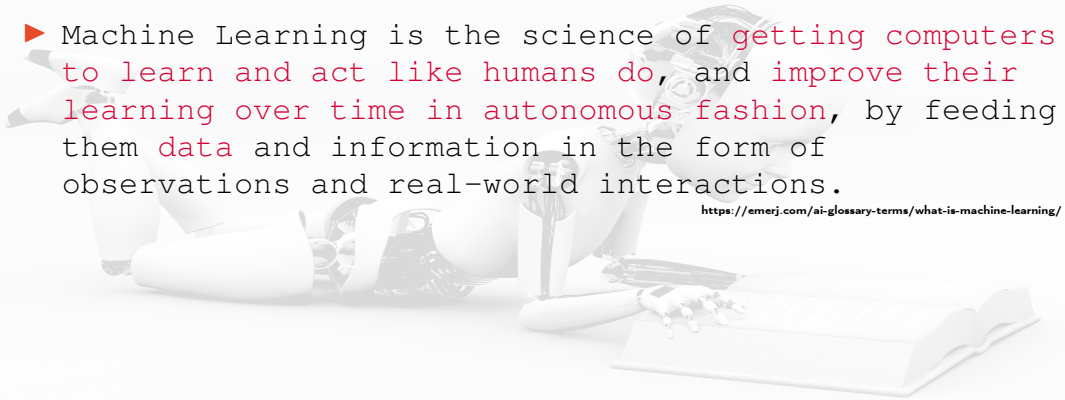

What is Machine Learning?

Definitions out there...

- Machine Learning is the science of getting computers to learn and act like humans do, and improve their learning over time in autonomous fashion, by feeding them data and information in the form of observations and real-world interactions.

https://emerj.com/ai-glossary-terms/what-is-machine-learning/

- Machine Learning at its most basic is the practice of using algorithms to parse data, learn from it, and then make a determination or prediction about something in the world.

https://blogs.nvidia.com/blog/2016/07/29/whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/

What is Machine Learning?

Definitions out there...

- Machine learning algorithms can figure out how to perform important tasks by generalizing from examples.

http://homes.cs.washington.edu/ pedrod/papers/cacm12.pdf

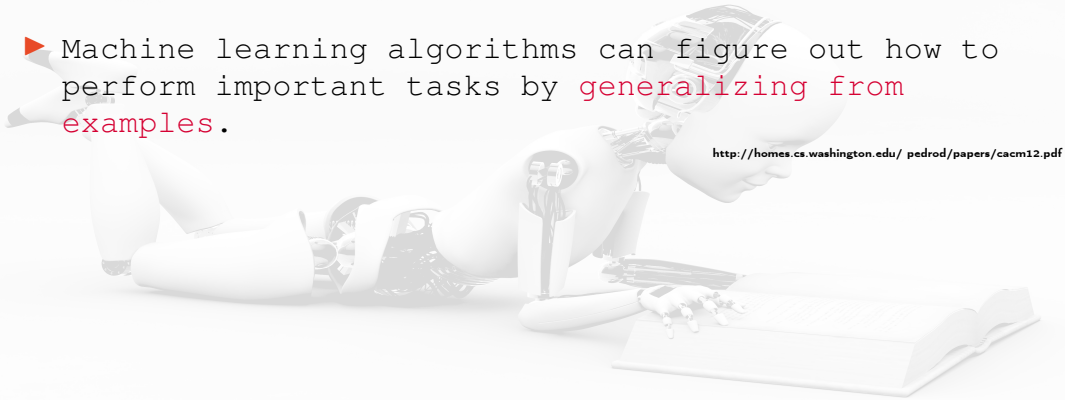

What is Machine Learning?

Definitions out there...

- Machine learning algorithms can figure out how to perform important tasks by generalizing from examples.

http://homes.cs.washington.edu/ pedrod/papers/cacm12.pdf

- The field of Machine Learning seeks to answer the question "How can we build computer systems that automatically improve with experience, and what are the fundamental laws that govern all learning processes?"

http://www.cs.cmu.edu/ tom/pubs/MachineLearning.pdf

What is Machine Learning?

Let's ask the machine...

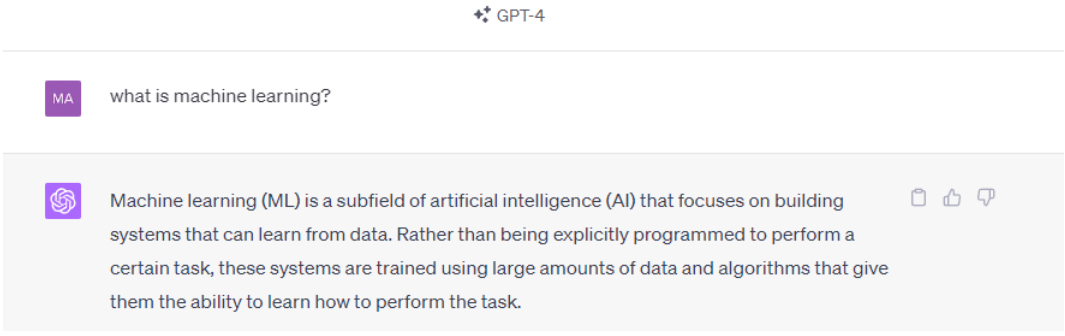

What is Machine Learning?

Traditional Programming

What is Machine Learning?

ML Programming

Machine Learning is the field of study that gives computers the ability to learn without being explicitly programmed
A.L. Samuel (1959)

What is Machine Learning?

ML Programming: "Harry Potter's approach"

Machine Learning is the field of study that gives computers the ability to learn without being explicitly programmed
A.L. Samuel (1959)

What is Machine Learning?

ML Programming: a more realistic approach

Programmer

Create a code based a
mathematical/statistical model
Defined by an ***expert****.*

Outputs

What is Machine Learning?

An Example

Is this a valid email address: marcelom@illinois.edu?

Traditional Programming

IF

upper/lowercase letters or digits

AND

"@" followed by a valid domain

AND NOT

special characters "#@*()..."

AND

...

THEN

valid

ELSE

invalid

Machine Learning

Input: (Big) data on valid/invalid email addresses

Y/N	Address
0	john@williams
1	xyz@greatcompany.com
0	mister.paul@GOT†.rxs
...	...

Computer: Statistical model where

$$Y/N = f\,(\text{input}\,|\,\text{parameters})$$

Output: Estimated probability of valid email

What is Machine Learning?

- Well, what do we want to learn/do?

What is Machine Learning?

- Well, what do we want to learn/do?
  - To predict or classify data

What is Machine Learning?

- Well, what do we want to learn/do?
  - To predict or classify data
  - To make causal predictions (causal inference)

What is Machine Learning?

- Well, what do we want to learn/do?
  - To predict or classify data
  - To make causal predictions (causal inference)
  - To group objects (uncover cluster structures)

What is Machine Learning?

- Well, what do we want to learn/do?
  - To predict or classify data
  - To make causal predictions (causal inference)
  - To group objects (uncover cluster structures)
  - To learn structural parameters (economics, epidemiology, ...)

What is Machine Learning?

- Well, what do we want to learn/do?
  - To predict or classify data
  - To make causal predictions (causal inference)
  - To group objects (uncover cluster structures)
  - To learn structural parameters (economics, epidemiology, ...)
  - To generate new data

What is Machine Learning?

- Well, what do we want to learn/do?
  - To predict or classify data
  - To make causal predictions (causal inference)
  - To group objects (uncover cluster structures)
  - To learn structural parameters (economics, epidemiology, ...)
  - To generate new data
  - ...

What is Machine Learning?

- Well, what do we want to learn/do?
  - To predict or classify data
  - To make causal predictions (causal inference)
  - To group objects (uncover cluster structures)
  - To learn structural parameters (economics, epidemiology, ...)
  - To generate new data
  - ...
- From what do we want to learn?

What is Machine Learning?

- Well, what do we want to learn/do?
  - To predict or classify data
  - To make causal predictions (causal inference)
  - To group objects (uncover cluster structures)
  - To learn structural parameters (economics, epidemiology, ...)
  - To generate new data
  - ...
- From what do we want to learn?
  - From data!

What is Machine Learning (ML)?

My view

- Automated computer **algorithms/methods + statistical models** to "learn" (discover) hidden patterns from data.

What is Machine Learning (ML)?

My view

- Automated computer **algorithms/methods + statistical models** to "learn" (discover) hidden patterns from data.
- Interpretation/explanation of the model/method is now an essential component of Machine Learning solutions.

What is Machine Learning (ML)?

My view

- Automated computer **algorithms/methods + statistical models** to "learn" (discover) hidden patterns from data.
- Interpretation/explanation of the model/method is now an essential component of Machine Learning solutions.

Ingredient List

Data + empirical question (what we want to learn) +

(Statistical) Model + loss function

+ optimization algorithms + evaluation/intepretation tools

Machine Learning Models/Methods

Different types of ML models/methods

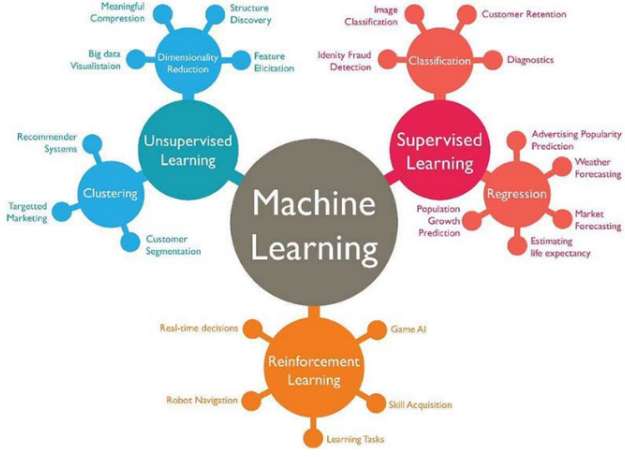

Source: https://www.datasciencecentral.com/profiles/blogs/types-of-machine-learning-algorithms-in-one-picture

Evolution of ML Models/Methods

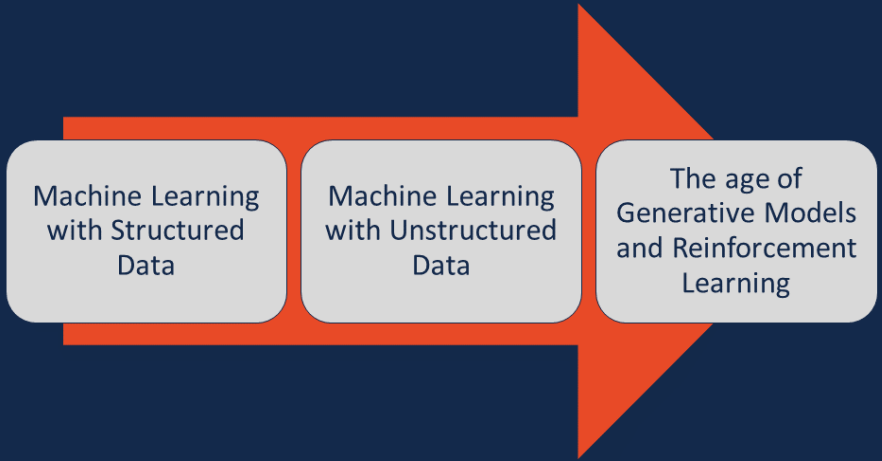

Evolution of ML Models/Methods

First Generation

- ML model previously trained with structured (tabular) data
- The ML model produces new outputs based on new data

Evolution of ML Models/Methods

Second Generation – The art of transforming text and images into numbers

- ML model previously trained with unstructured data (text, images, audio, ...)
- The ML model produces new outputs based on new data

Evolution of ML Models/Methods

Third Generation

- ML model previously trained with structured/unstructured data
- The ML model produces new outputs "upon request"

ML/AI in Economics and Finance

Machine Learning with Structured Data

Structured (numerical) data
Usually commonly available

LASSO & Friends
Tree-Based Models
Ensemble methods
Neural Network family
Model Combination
...

Large Scale Forecasting
(macro, sales, etc...)

Credit Models and
Fraud Detection

Portfolio Choices and
Investment Decisions

Marketing Strategies

Pricing

ML/AI in Economics and Finance

Machine Learning with Unstructured Data

Unstructured data available online

Text processing tools
Dictionaries
Count model (N-grams)
Topic Models (LDA)
Image processing tools

Structured (numerical) data
Not commonly available

LASSO & Friends
Tree-Based Models
Ensemble methods
Neural Network family
Model Combination
...

Large Scale Forecasting (macro, sales, etc...)

Credit Models and Fraud Detection

Portfolio Choices and Investment Decisions

Marketing Strategies

Pricing

ML/AI in Economics and Finance

The Age of AI and the Future Ahead…

Large Language Models
Summarize
Topic allocation
Sentiment Analysis

Image Processing Tools

Large Scale Forecasting (macro, sales, etc…)

Credit Models and Fraud Detection

Portfolio Choices and Investment Decisions

Marketing Strategies

Pricing

ML/AI in Economics and Finance

Products

News/Social Media Based Leading Indicators
Sentiment Analysis
Topic Models
...

Image processing
Revisiting technical analysis
Information from graphs

Document processing
Information from firms' reports
Social media scrapping
Central Bank communication

Large Scale Forecasting Models

Investment Strategies/Pricing

Credit Models/Fraud Detection

ML/AI in Economics and Finance

Products

News/Social Media Based Leading Indicators
Sentiment Analysis
Topic Models
...

Image processing
Revisiting technical analysis
Information from graphs

Document processing
Information from firms' reports
Social media scrapping
Central Bank communication

Large Scale Forecasting Models

Investment Strategies/Pricing

Credit Models/Fraud Detection

What's next?

See you next class...