2695 Introduction to Machine Learning
Masters Program in Economics, Finance and Management

NOVA SCHOOL OF BUSINESS & ECONOMICS

1a  Video Lecture

INTRODUCTION TO MACHINE LEARNING

Is Machine Learning new?

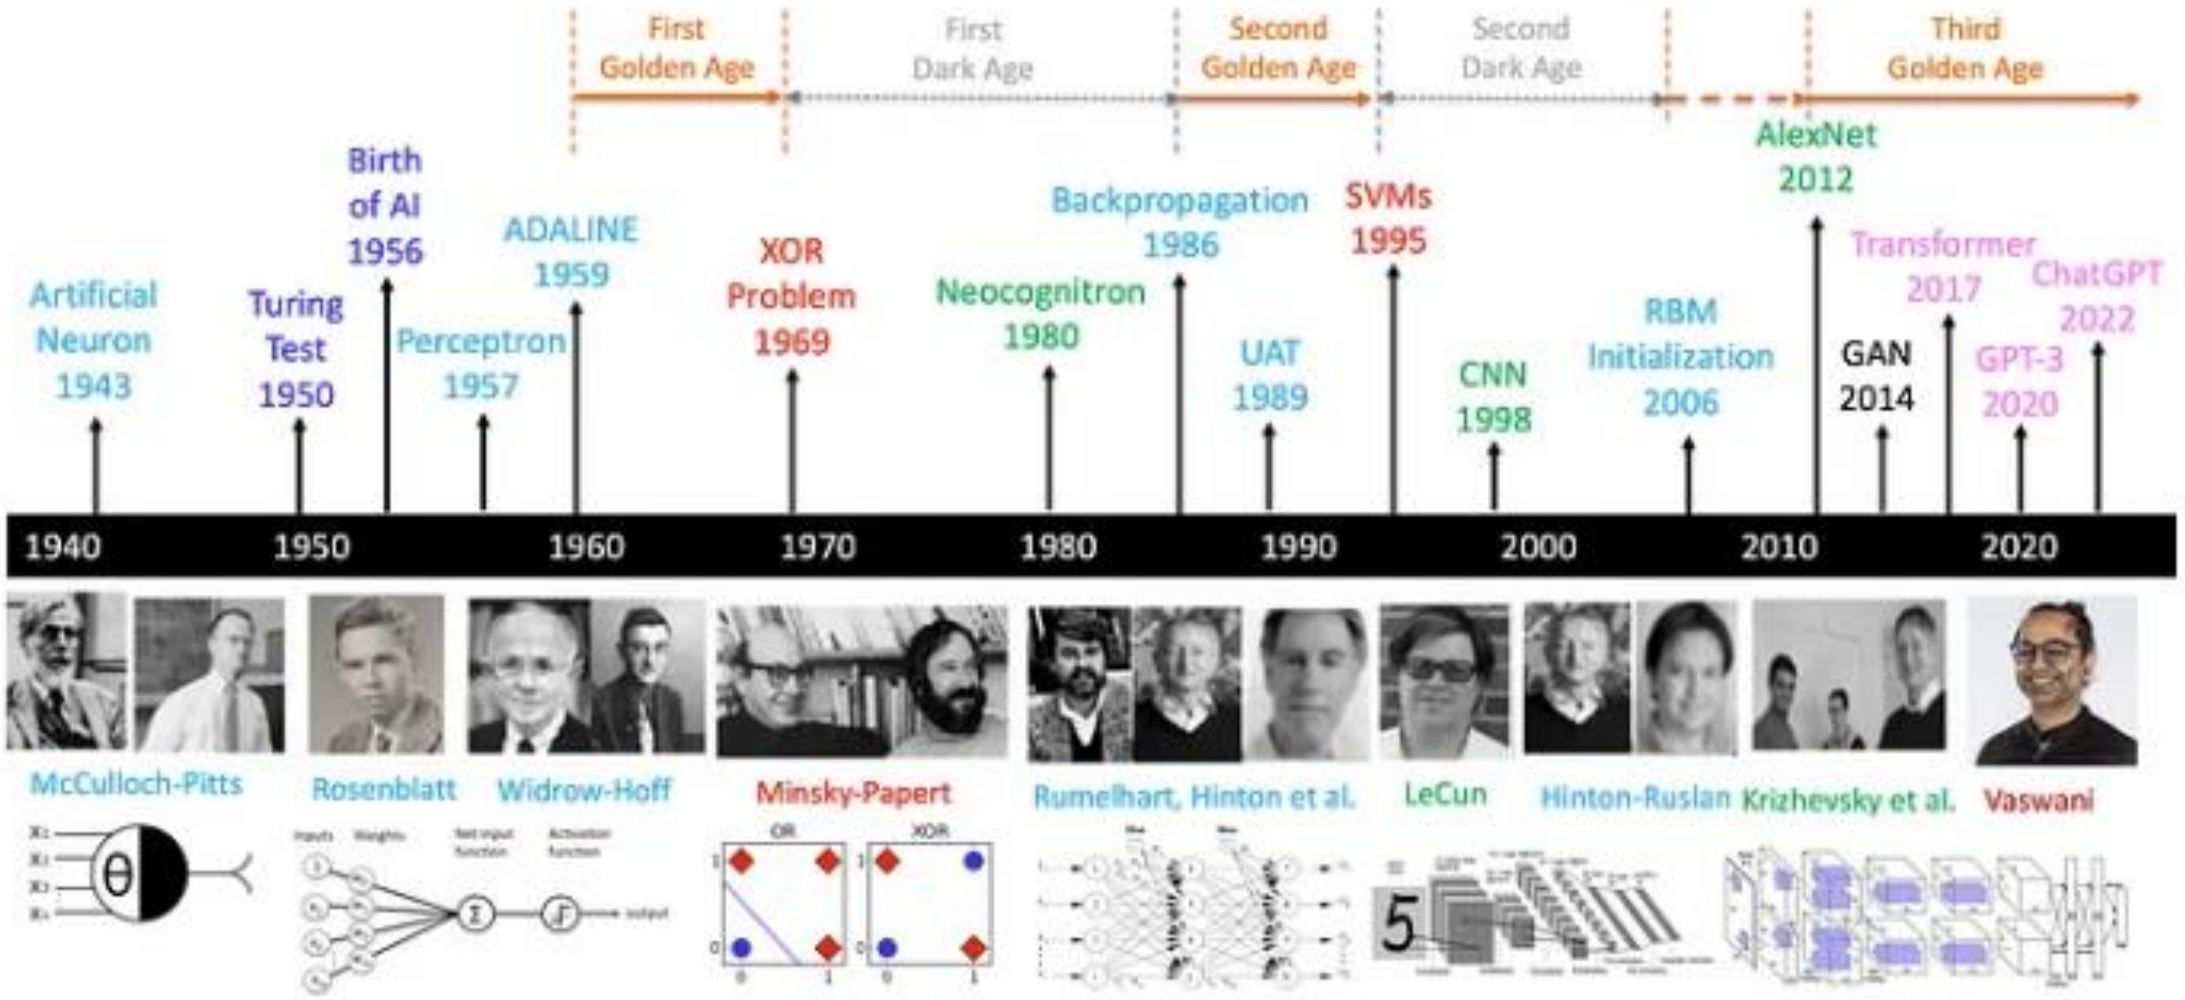

A Brief History of AI with Deep Learning, LM Po

How AI is changing the world today

- **Economics**: macroeconomics forecasting, market analysis
- **Finance**: fraud detection, algorithmic trading
- **Retail**: recommendation engines, demand forecasting
- **Transportation**: autonomous vehicles, traffic management
- **Energy**: optimizing wind and solar power generation, smart home systems

Future employment in the era of AI

World Economic Forum's Future of Jobs Report 2025: Fastest growing and declining jobs by 2030

	Top fastest growing jobs		Top fastest declining jobs
1	Big data specialists	1	Postal service clerks
2	FinTech engineers	2	Bank tellers and related clerks
3	AI and machine learning specialists	3	Data entry clerks
4	Software and applications developers	4	Cashiers and ticket clerks
5	Security management specialists	5	Administrative assistants and executive secretaries
6	Data warehousing specialists	6	Printing and related trades workers
7	Autonomous and electric vehicle specialists	7	Accounting, bookkeeping and payroll clerks
8	UI and UX designers	8	Material-recording and stock-keeping clerks
9	Light truck or delivery services drivers	9	Transportation attendants and conductors
10	Internet of things specialists	10	Door-to-door sales workers, news and street vendors, and related workers
11	Data analysts and scientists	11	Graphic designers
12	Environmental engineers	12	Claims adjusters, examiners and investigators
13	Information security analysts	13	Legal officials
14	DevOps engineers	14	Legal secretaries
15	Renewable energy engineers	15	Telemarketers

https://www.weforum.org/stories/2025/01/future-of-jobs-report-2025-the-fastest-growing-and-declining-jobs/

NOVA SCHOOL OF BUSINESS & ECONOMICS

1b Video Lecture

AI vs DATA SCIENCE vs MACHINE LEARNING

What is Artificial Intelligence

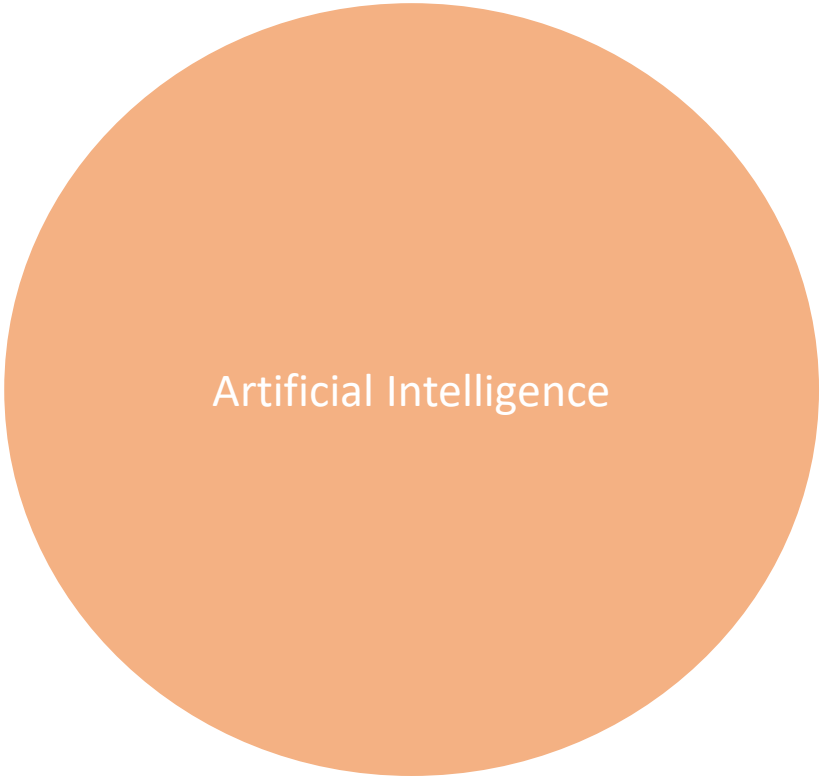

Ability of a computer to perform tasks commonly associated with intelligent being

What is Machine Learning

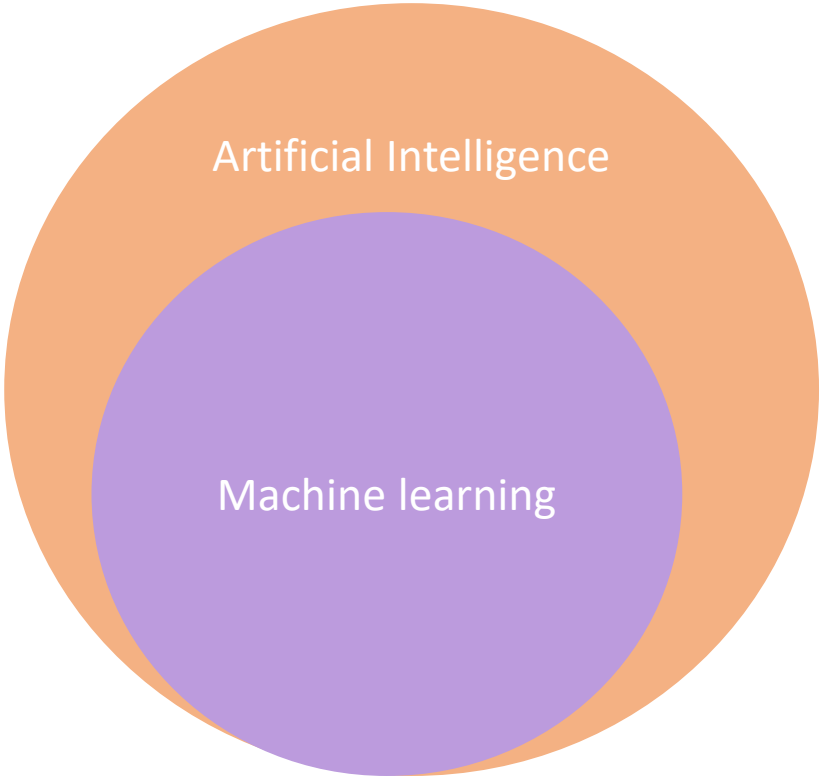

Science and art of programming computers so they can *learn from data*.

Study of algorithms that improve their performance at some task with experience.

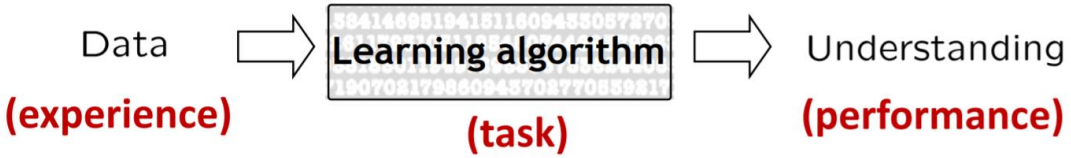

What is Deep Learning

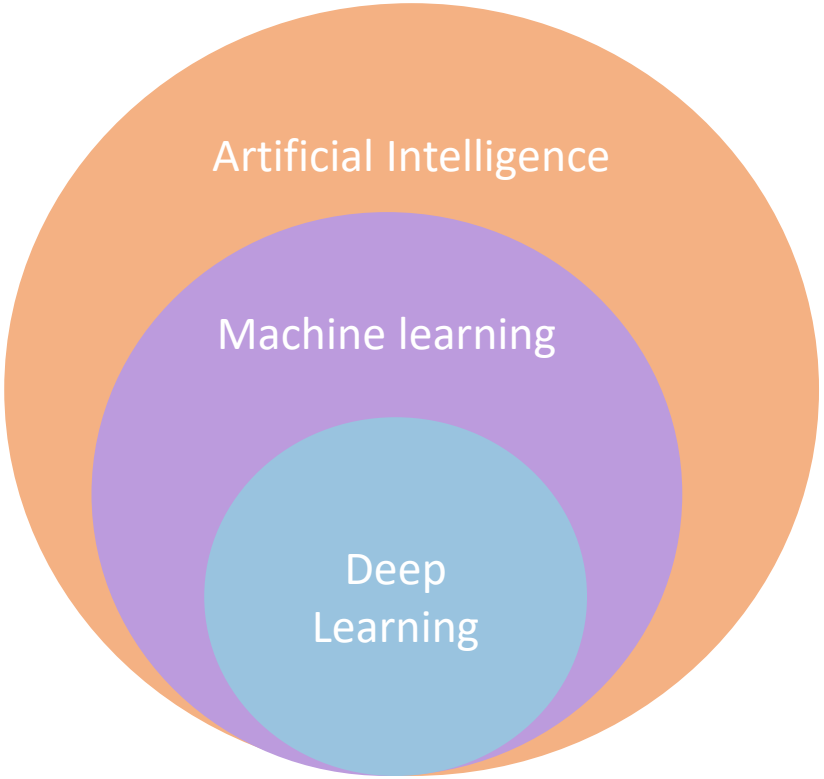

In comparison to ML, Deep Learning:

- can solve more complex problems
- is more difficult to implement
- requires more time to train the model
- requires more computational resources.

What is Data Science

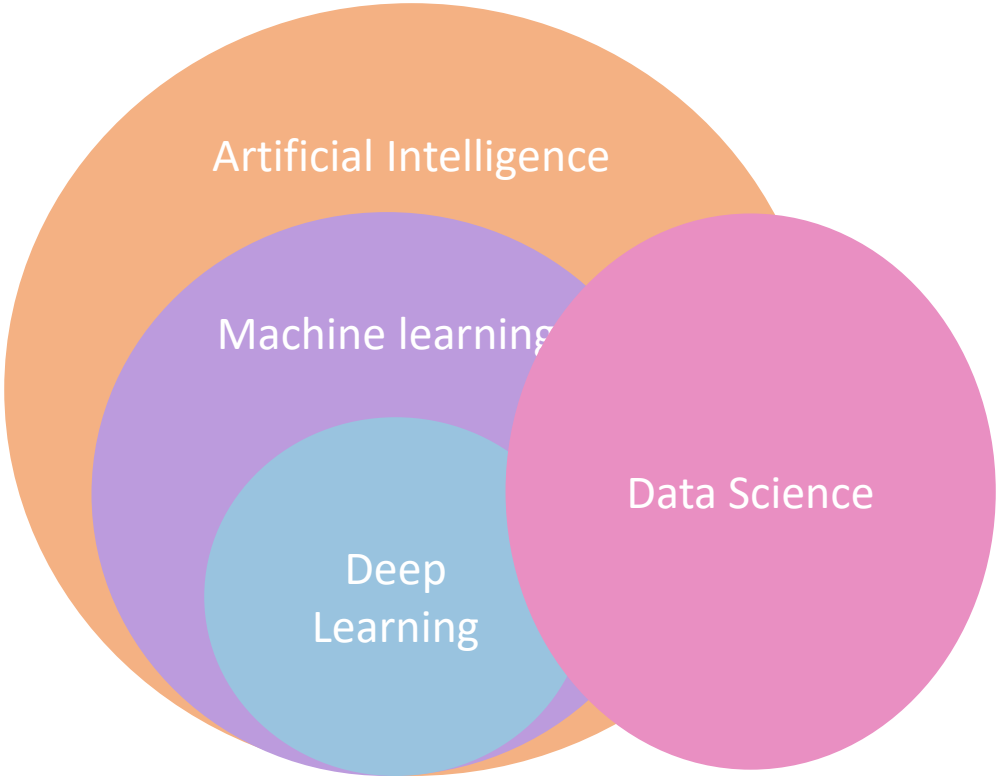

Set of fundamental principles that guide the extraction of knowledge from data.

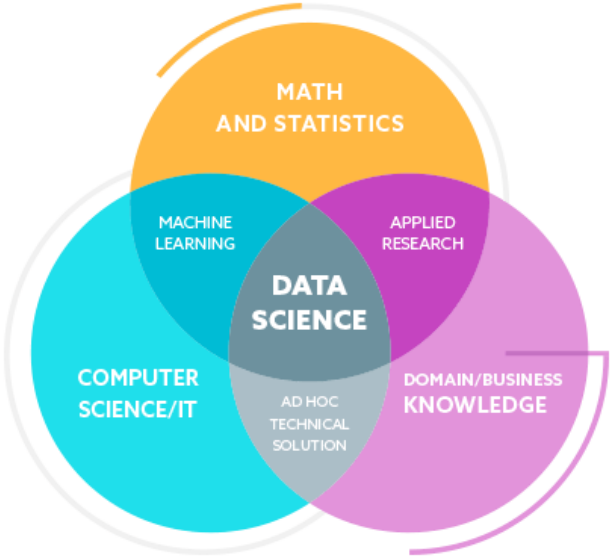

Key Takeaways: Difference between Data Science, Machine Learning and AI

The fields do have a great deal of overlap, but they are **not interchangeable.**

1. Data science produces actionable **insights.**
2. Machine learning produces **predictions.**
3. Artificial intelligence produces **actions.**

NOVA SCHOOL OF BUSINESS & ECONOMICS

1c Video Lecture

HOW AND WHEN TO USE ML FOR BUSINESS

Data mining process

Cross Industry Standard Process for Data Mining (CRISP-DM; Shearer, 2000)

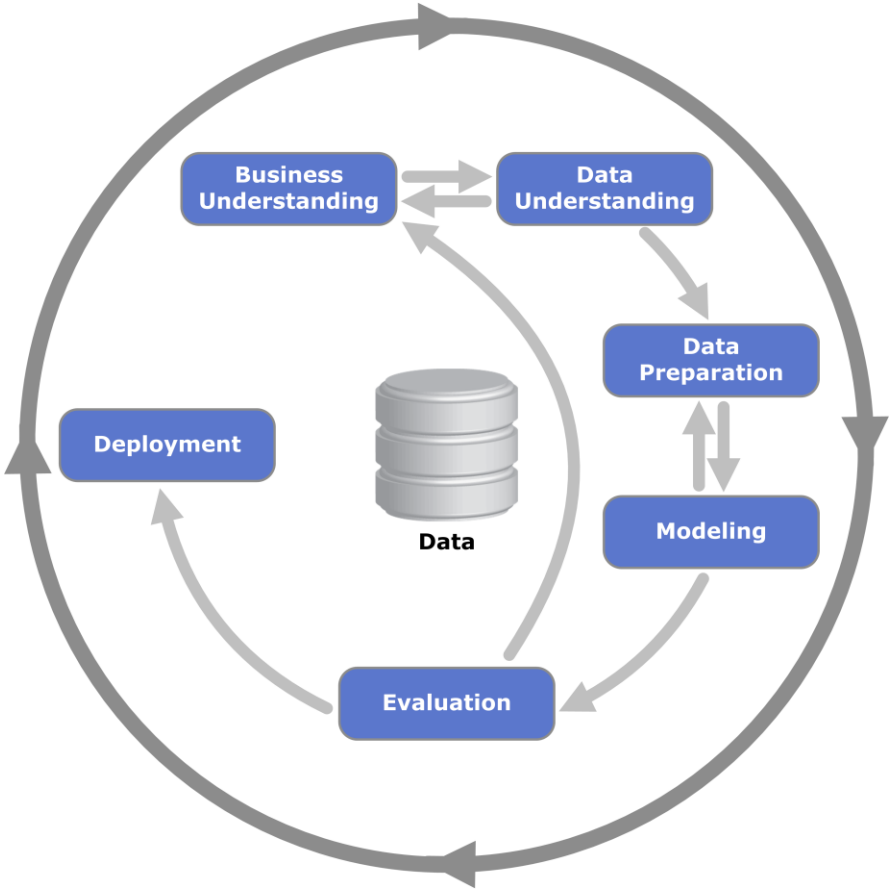

CRISP-DM slides content credit: Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking, Foster Provost and Tom Fawcett

Business understanding

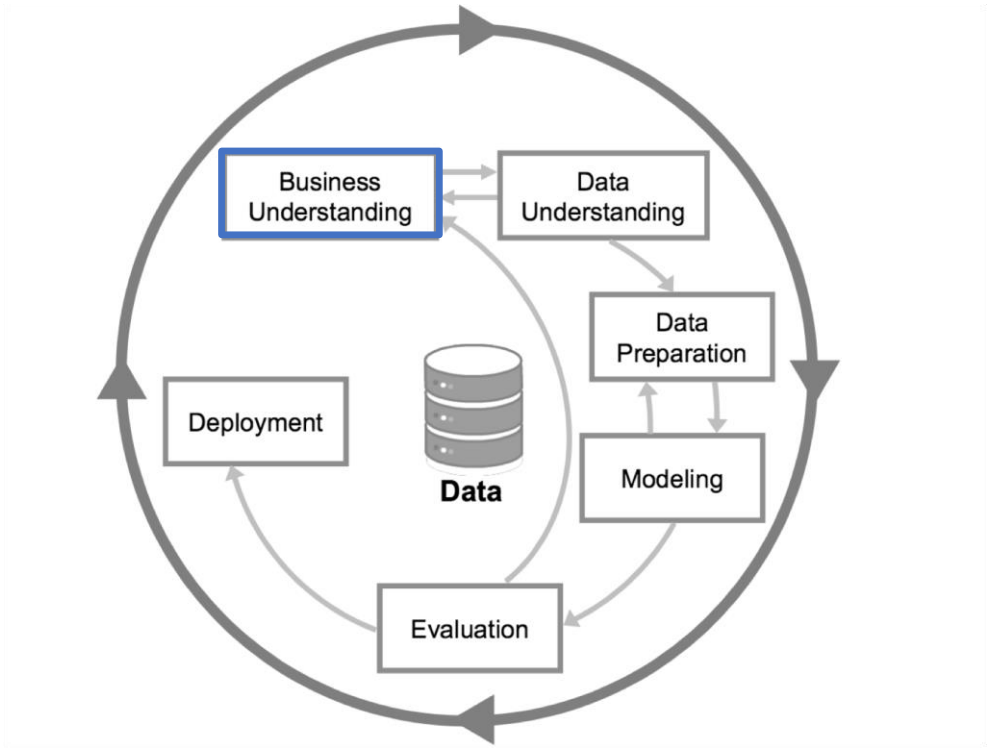

- Identify our business goals
- Assess our situation
- Define our data science goals
- Produce our project plan

CRISP-DM slides content credit: Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking, Foster Provost and Tom Fawcett

Data understanding

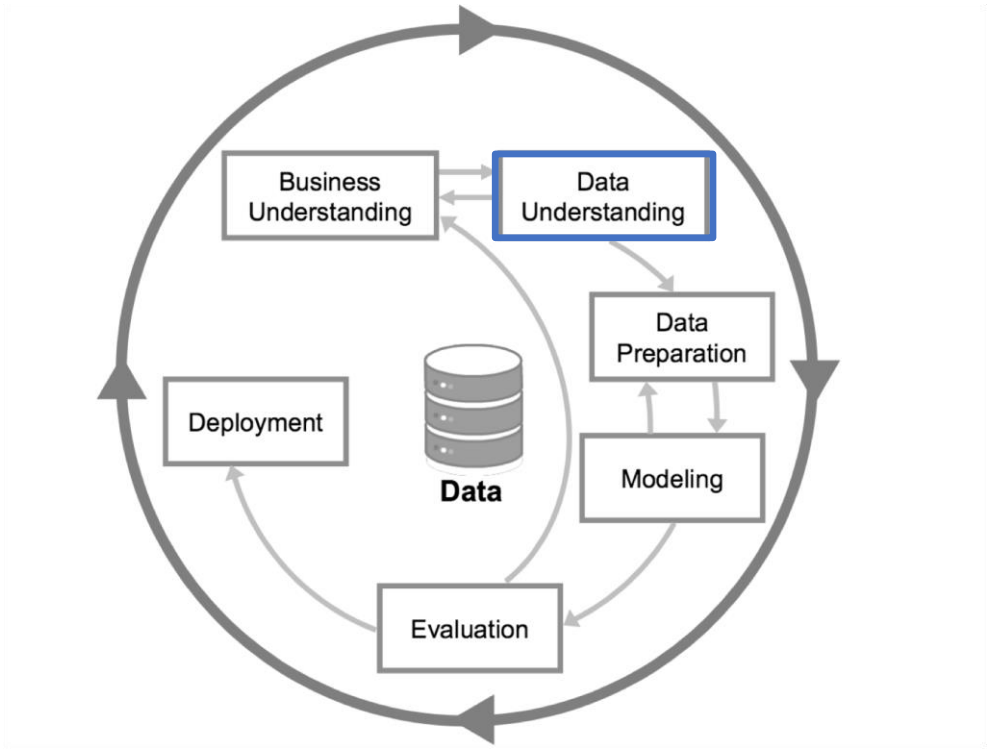

- Collect initial data
- Describe data
- Explore data
- Verify data quality

CRISP-DM slides content credit: Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking, Foster Provost and Tom Fawcett

Data preparation

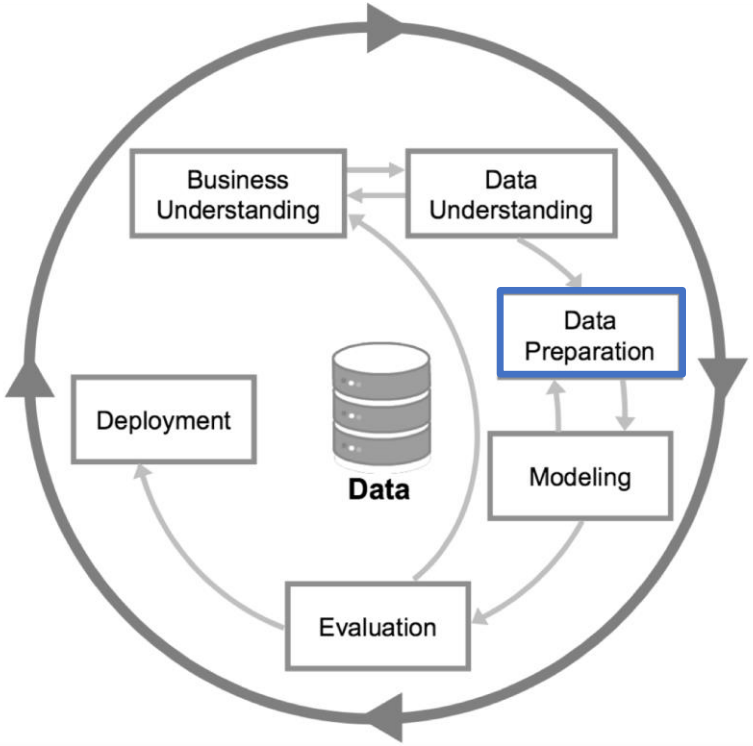

- Data preparation/curation/preprocessing
  - Invalid values
  - Formats
  - Uniqueness
  - Missing values
  - Misspellings

CRISP-DM slides content credit: Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking, Foster Provost and Tom Fawcett

Modeling

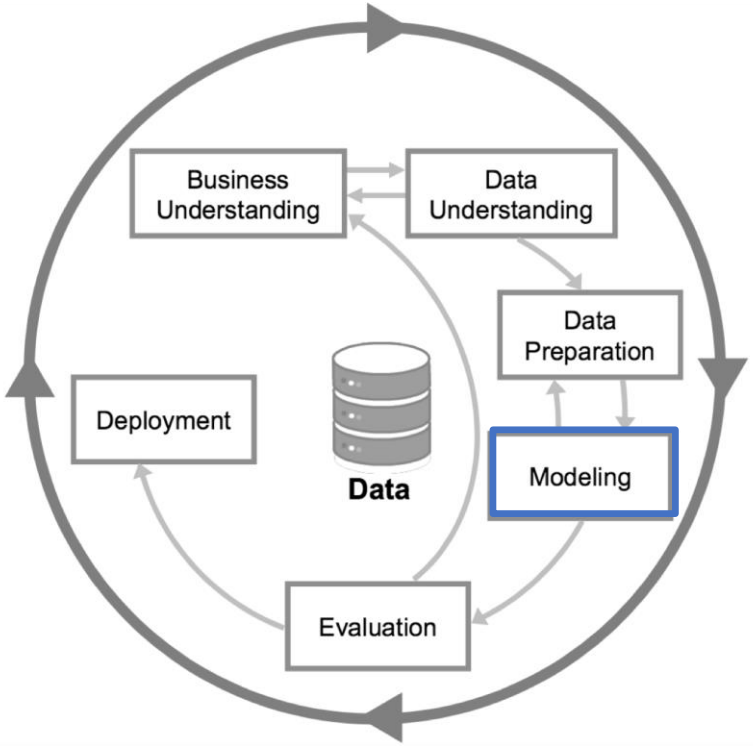

Build a model that captures regularities of data.

CRISP-DM slides content credit: Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking, Foster Provost and Tom Fawcett

Model evaluation

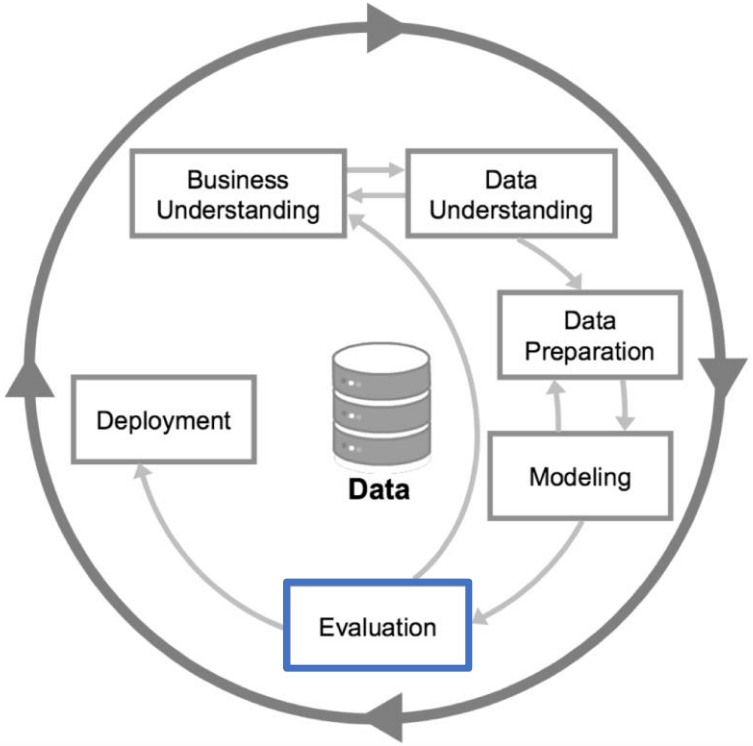

- Evaluate performance
- Make sure that the model satisfies initial business goals
- Interpret results

CRISP-DM slides content credit: Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking, Foster Provost and Tom Fawcett

Model deployment

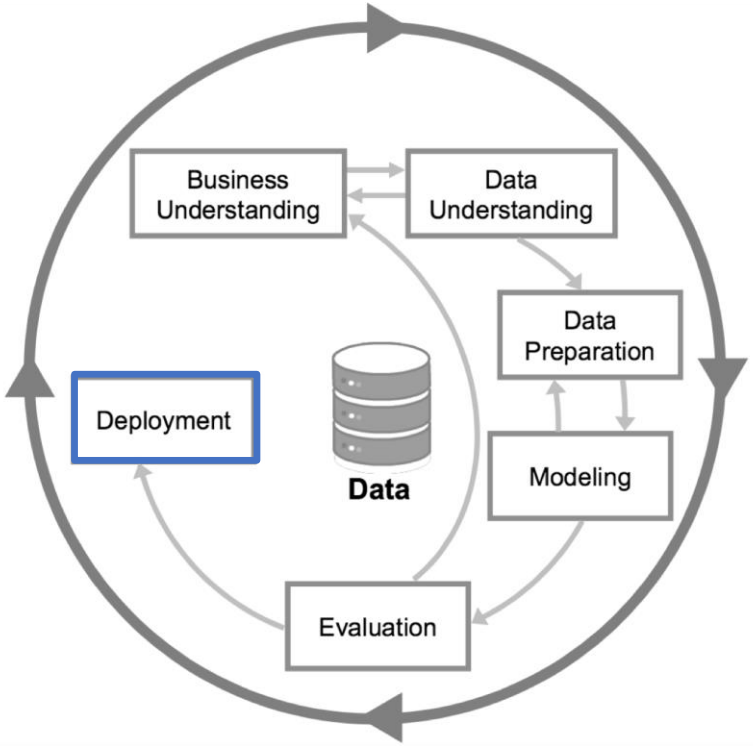

We might need to solve a business problem:

- Once (Specific campaign)
- Regularly (Monthly campaign of contacting customers). Here, we need to integrate the machine learning model into a production environment.

CRISP-DM slides content credit: Data Science for Business: What You Need to Know about Data Mining and Data-Analytic Thinking, Foster Provost and Tom Fawcett

After model deployment

MLOps is a set of practices that aims to deploy and maintain machine learning models in production reliably and efficiently.

https://ml-ops.org/content/mlops-principles

When should you use ML

Can you formulate your problem clearly?

Do you need to automate your task?

Do you have sufficient data?

Does your problem have a regular pattern?

How do you define success?

When should you use ML: example of sentiment analysis

Can you formulate your problem clearly?

Given a customer review, infer its sentiment:

- Input: customer review text
- Output: positive, negative, neutral

Do you need to automate your task?

Too many review to read.

Do you have sufficient data?

Millions of customer reviews and ratings from the Web.

Does your problem have a regular pattern?

High rated reviews more positive words (*good, awesome)*
Low rated reviews more negative words (*bad, disappointed)*

How do you define success?

Percentage of reviews with correctly inferred sentiment.

More examples of when to use ML

Hand-written rules and equations are too complex—as in face recognition and speech recognition.

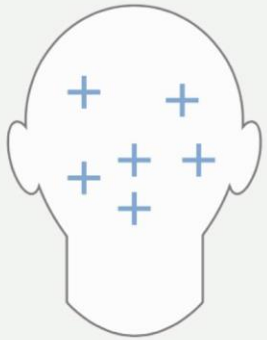

The rules of a task are constantly changing—as in fraud detection from transaction records.

The nature of the data keeps changing, and the program needs to adapt—as in automated trading, energy demand forecasting, and predicting shopping trends.