Machine Learning in Medicine

Tran Giang Son, tran-giang.son@usth.edu.vn

ICT Department, USTH

Review

AI, ML, DL, DS

- Whaaaaaaaaaat?

AI, ML, DL, DS

- Whaaaaaaaaaat?

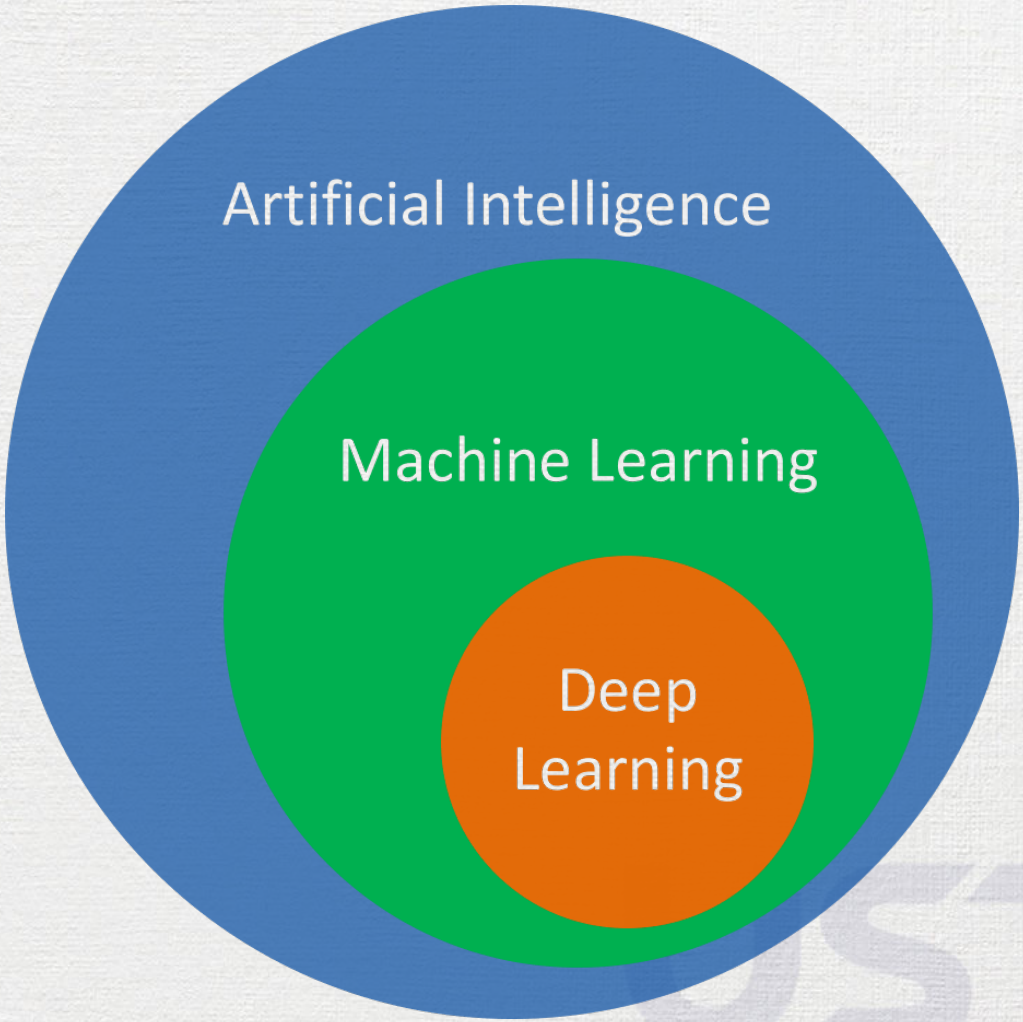

Artificial Intelligence

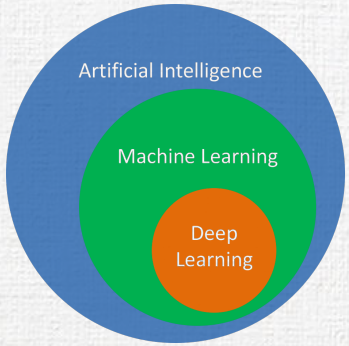

- The effort to automate intellectual tasks normally performed by humans
- Many more approaches without learning
  - Symbolic AI, Hardcoded rules
  - Expert systems
- Simply too many rules...

Machine Learning

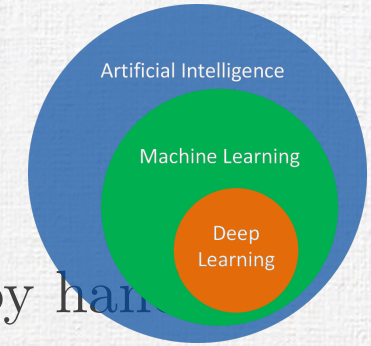

- AI: Programmers crafting data-processing rules by hand
- A computer learn these rules by looking at data?
  - *Trained*
  - Not explicitly programmed
  - Statistical rules
- Kinda the most successful subfield of AI
  - More data
  - Bigger, faster hardware

Machine Learning

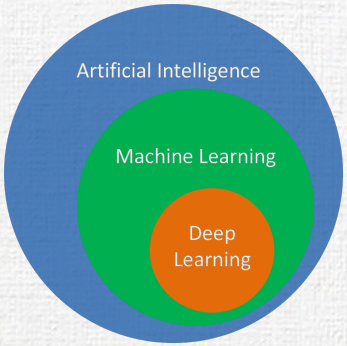

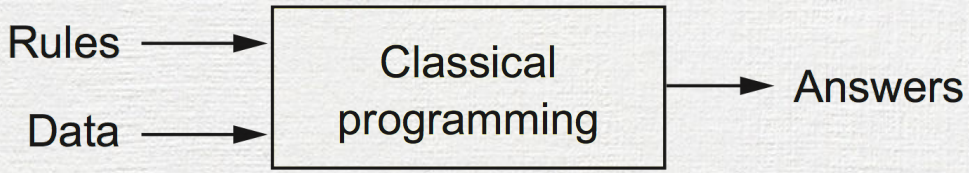

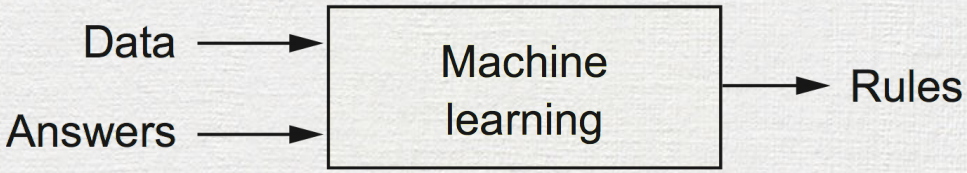

Machine Learning

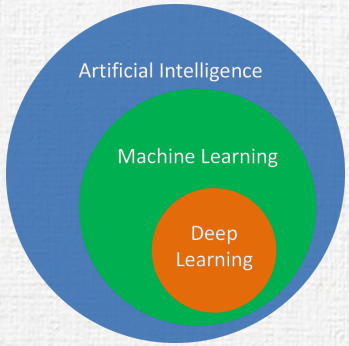

- Input data
- Examples of expected output
  - Also called
- Evaluating if algorithm is good

Machine Learning: Input

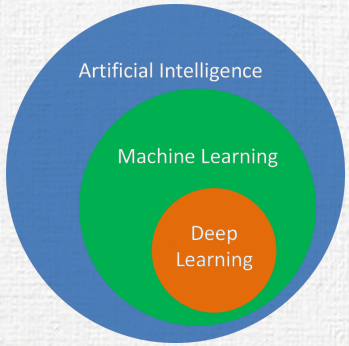

- Data abstraction
  - Pixels
  - Edges
  - Shapes
  - Human
  - Crowd

Machine Learning: Expected output

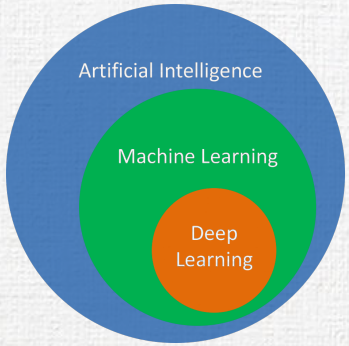

- Expected output
  - Pixels: red, green, blue
  - Edges: horizontal, vertical, diagonal
  - Shapes: rectangle, circle, oval
  - Human: head, hand, leg, body
  - Crowd: male, female, kid, toddler

Machine Learning: evaluation

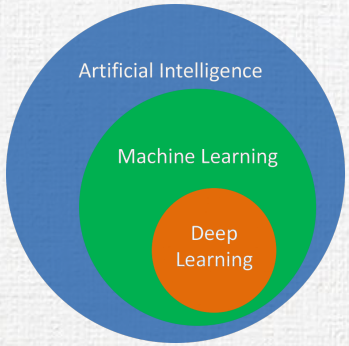

- Input: a photo

Machine Learning: evaluation

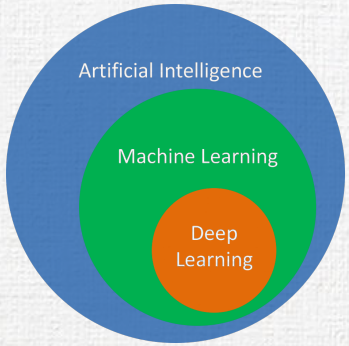

- Input: a photo
- Expected output: a dog

Machine Learning: evaluation

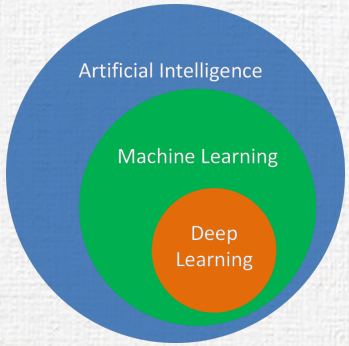

- Input: a photo
- Expected output: a dog
- Algorithm:

Machine Learning: evaluation

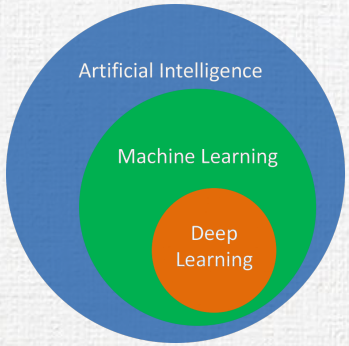

- Input: a photo
- Expected output: a dog
- Algorithm:
  - "I guess this is a cat"

Machine Learning: evaluation

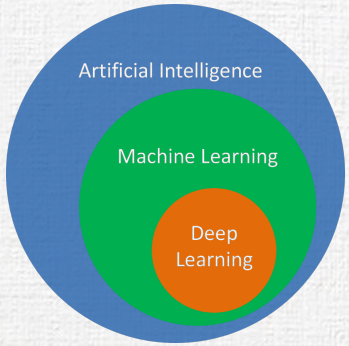

- Input: a photo
- Expected output: a dog
- Algorithm:
  - "I guess this is a cat"
  - Looks at the output

Machine Learning: evaluation

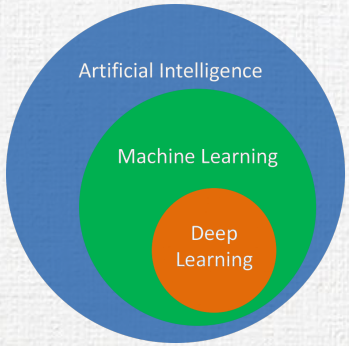

- Input: a photo
- Expected output: a dog
- Algorithm:
  - "I guess this is a cat"
  - Looks at the output
    - "Err... I think it should have been a cat"

Machine Learning: evaluation

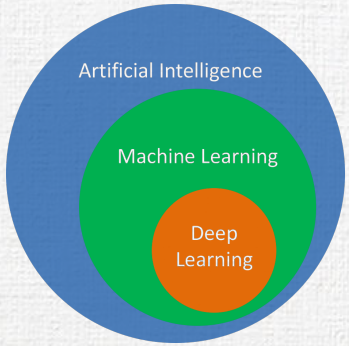

- Input: a photo
- Expected output: a dog
- Algorithm:
  - "I guess this is a cat"
  - Looks at the output
    - "Err... I think it should have been a cat"
    - "But they look like a cat"

Machine Learning: evaluation

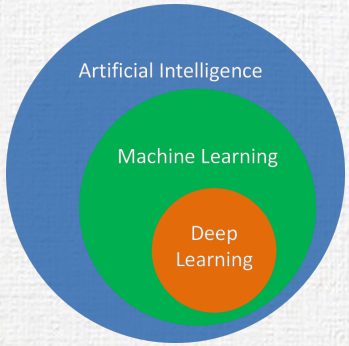

- Input: a photo
- Expected output: a dog
- Algorithm:
  - "I guess this is a cat"
  - Looks at the output
    - "Err... I think it should have been a cat"
    - "But they look like a cat"
    - "What are the differences?"

Machine Learning: evaluation

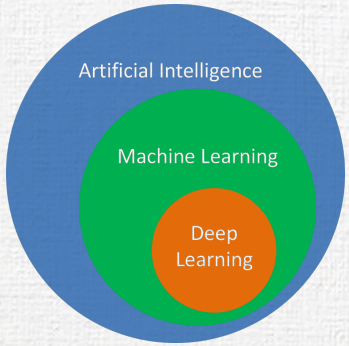

- Input: a photo
- Expected output: a dog
- Algorithm:
  - "I guess this is a cat"
  - Looks at the output
    - "Err... I think it should have been a cat"
    - "But they look like a cat"
    - "What are the differences?"
- Differences: evaluation between prediction and groundtruth

Machine Learning

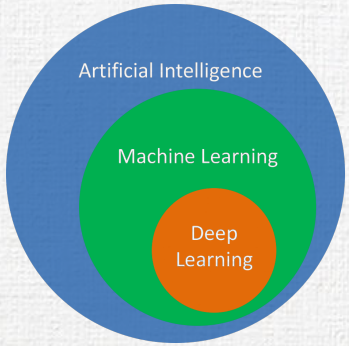

- Supervised learning
  - This is a dog, this is a cat
  - Does this photo look like a dog or a cat?
- Unsupervised learning
  - This one looks alike this one. Group 1.
  - That one looks alike another one. Group 2.

Machine Learning

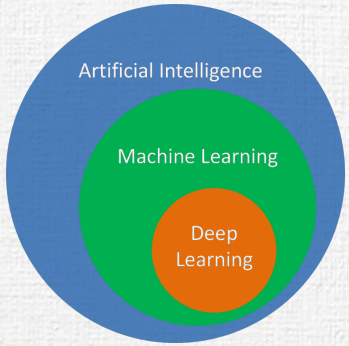

- Problems: classification, clustering regression, association rules, reinforcement learning, feature learning,...
  - Classification: decision tree, random forest, k-NN, linear regression, naive bayes, multi layer perceptron, relevance vector machine, support vector machine
  - Clustering: hierachical, k-means, mean-shift
  - ...

Machine Learning vs Data Mining

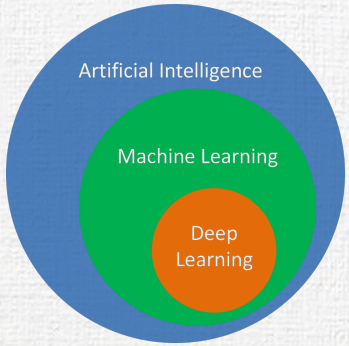

- Machine Learning: prediction based on known properties
- Data Mining: discovery of unknown properties

Deep Learning

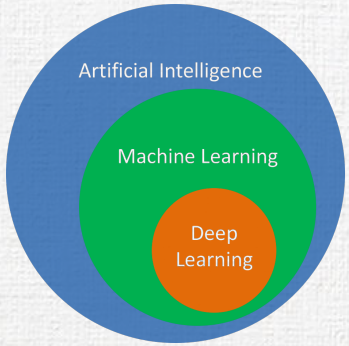

- A specific subfield of machine learning
- Learning data representation
  - Using **layers** of abstraction
  - Not deeper understanding

Deep Learning

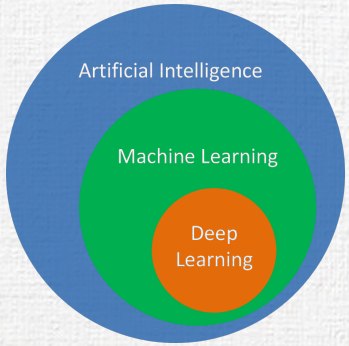

- A specific subfield of machine learning
- Learning data representation
  - Using **layers** of abstraction
  - Not deeper understanding
- Deeeeeeeeeeeeeeeeep in layers
  - Tens
  - Hundreds

Data Science

Data Science

- A science about data ☺

Data Science

- A science about data ☺
- Usage of deep learning, machine learning, artificial intelligence for a specific domain

Data Science

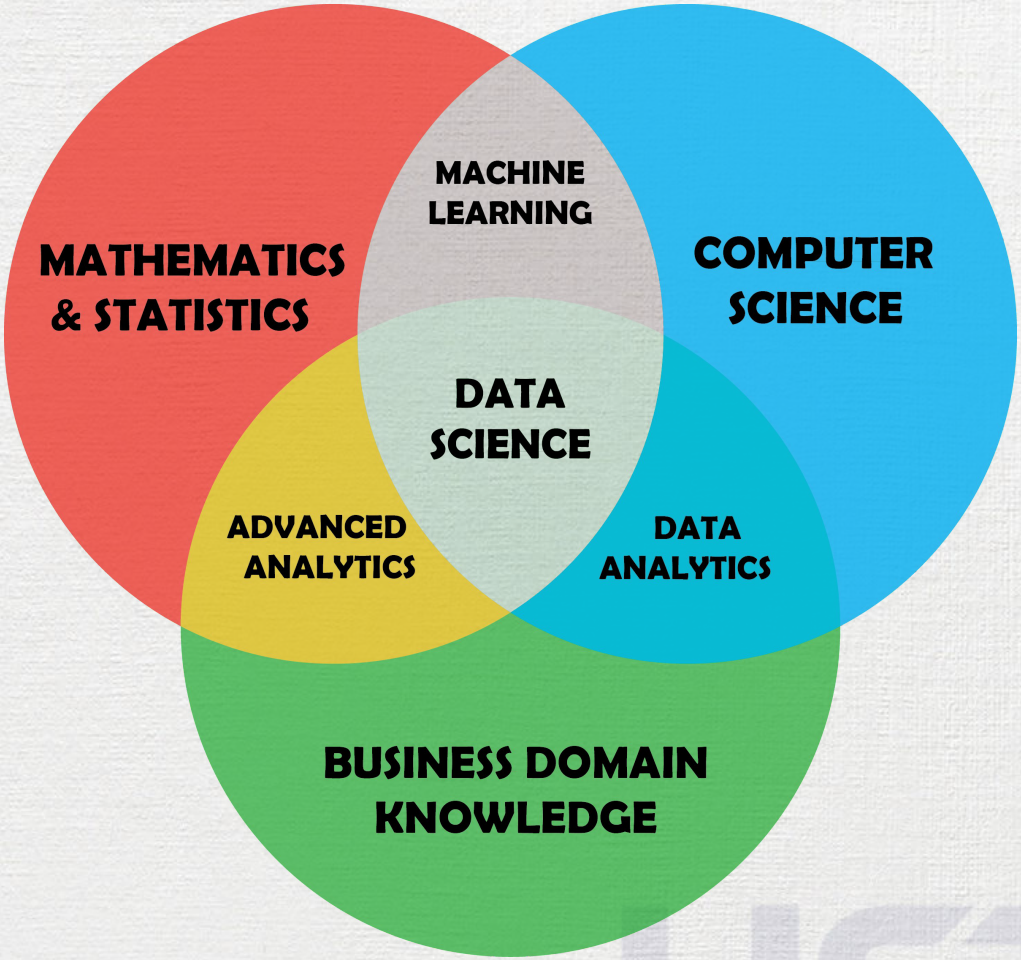

Medicine

Medicine

the science or practice of the diagnosis, treatment, and prevention of disease

Medicine

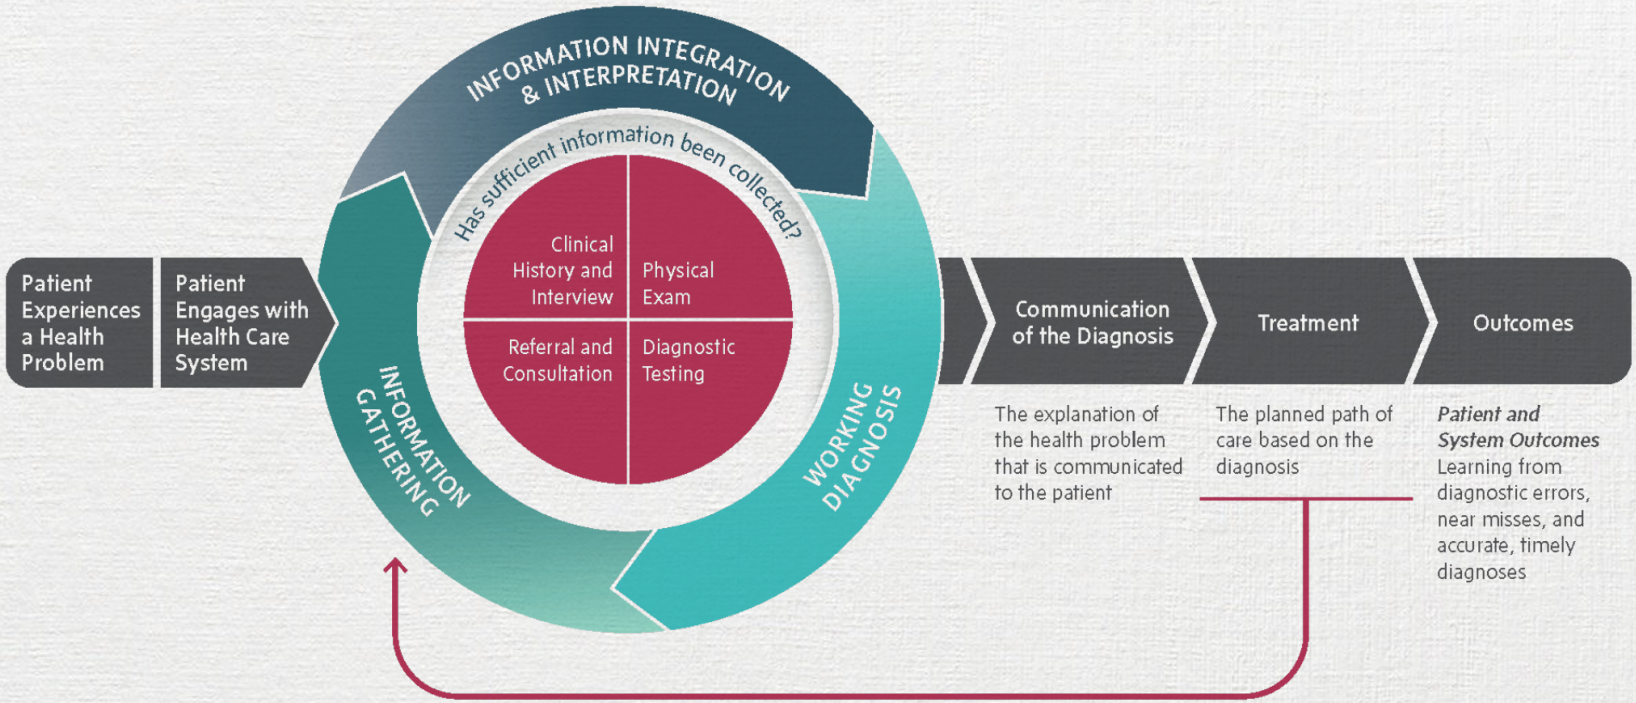

Pipeline

Medicine

- Information gathering
  - Clinical history
  - Interview
  - Physical exam
  - Diagnostic testing
- Information integration and interpretation
  - Comparing with knowledge and experience
  - Generating a working diagnosis

Medicine

- Determining a working diagnosis
  - Communicating the diagnosis to the patient
  - Degree of certainty
  - Next steps
- Feedback and follow-up
  - Monitoring the patient's response to treatment
  - Updating the diagnosis if necessary.

Diagnosis

- Complex, patient-centered, collaborative activity
  - Information gathering
  - Clinical reasoning
  - Goal: determining a patient's health problem

Diagnosis

- Considerations in the diagnostic process
  - Diagnostic uncertainty
  - Time-consuming
  - Population trends
  - Diverse populations
  - Health disparities
  - Mental health

Treatment

- A very wide domain in medicine
- We don't do it :-)

Prevention

- A very wide domain in medicine
- We don't do it :-)

ML in Medicine

What

- Application of computational techniques
  - Allow computers to perform complex tasks related to healthcare
  - Mathematical models and algorithms
    - Learn from data
    - Make predictions or decisions

Why

- More computational power, storage, memory
- More medical data generation: health records, imaging, genomics
- Early diagnosis
- Overloaded hospitals
- Personalized health care
- Reduced costs

Applications

- Disease Diagnosis and Prediction
  - Analyzing patient data (e.g., symptoms, lab results)
  - Diagnose diseases (e.g., cancer, diabetes) and predict outcomes
  - Early detection improves treatment efficacy
- Medical Imaging
  - Enhancing radiology by automating image analysis
  - Assisting in detecting tumors, fractures, and abnormalities
  - Many modalities: X-rays, MRIs, and CT scans

Applications

- Drug Discovery and Development
  - Predicting drug interactions
  - Identifying potential candidates
  - Optimizing existing drugs
  - Accelerating research by analyzing molecular data

Applications

- Personalized Treatment
  - Preparing treatment plans based on individual patient
  - Genetics, lifestyle, and medical history
  - Reducing costs and improves efficiency

Applications

- Genomics and Precision Medicine:
  - Analyzing genomic data to identify genetic variants associated with diseases
  - Informing personalized treatment strategies
- Predictive Analytics
  - Predicting disease outbreaks, patient readmissions, and healthcare utilization
  - Preventive measures.

How?

- Common techniques
  - Regression models
  - Classification models
- Metrics