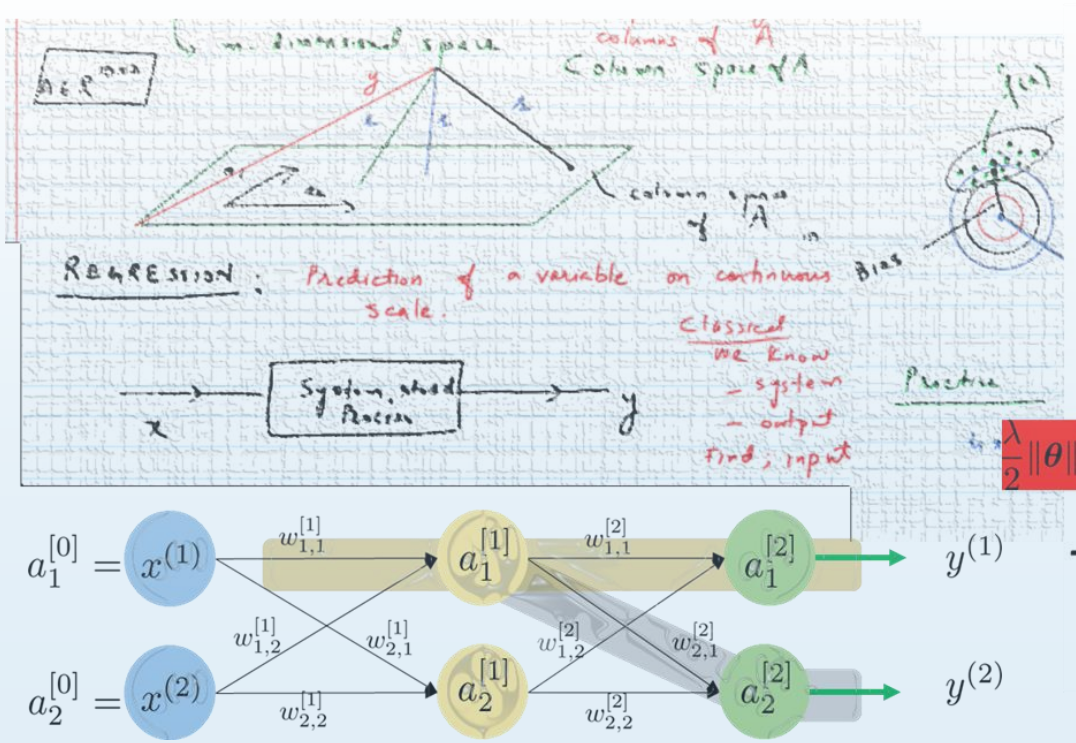


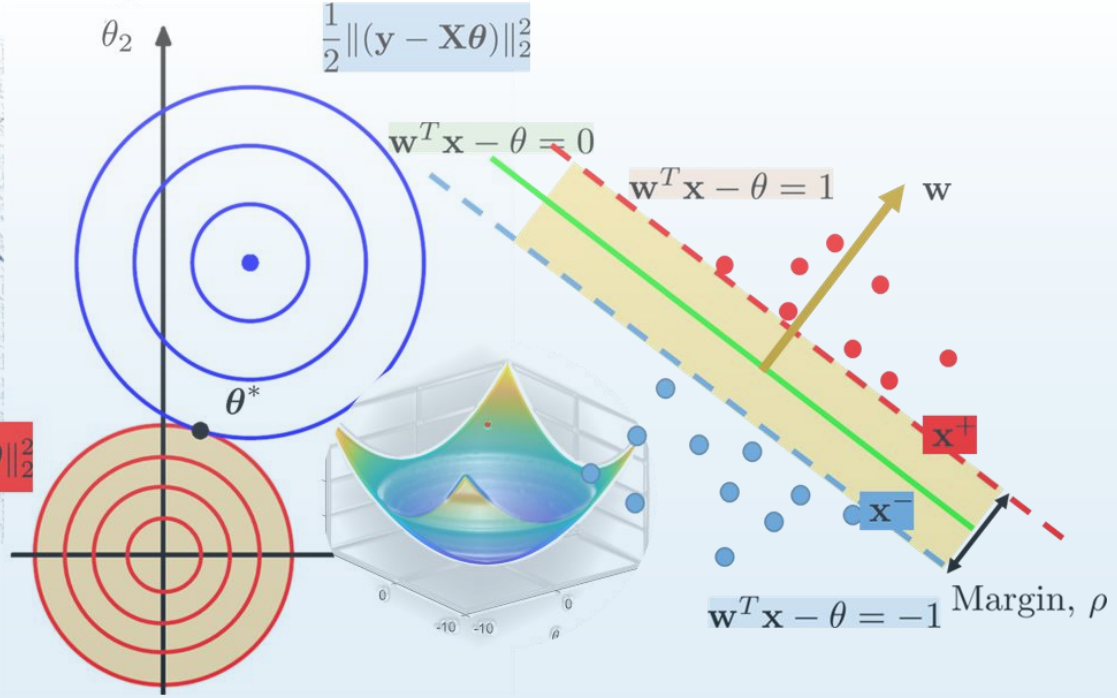


# AI-501 Mathematics for AI

## Overview

Zubair Khalid
School of Science and Engineering

https://www.zubairkhalid.org/ai501_2024.html

LUMS
A Not-for-Profit University

# Outline

- Course Overview, Logistics
- AI Overview
- Introduction to Machine Learning

# About us!

Salaar

Zubair

Umer

Ibrahim

Ayyan

Fatima

LUMS
A Not-for-Profit University

# About the Instructor

- Associate Professor, LUMS
- Post-doctorate – 2013-2015, Australian National University (ANU)
- PhD, Australian National University (ANU) - 2013

**Affiliations:**

- CITY – Centre for Urban Informatics, Technology and Policy (www.city.lums.edu.pk)
- Smart Data, Systems and Applications Lab (www.sdsa.lums.edu.pk)

**Collaborations:** Princeton, UCL, University of Edinburgh, EPFL, ANU, KAUST

**PhD Students:** 10 (7 graduated)

**Publications:** More than 85 (30 Transactions/Journals, 53 Conference proceedings)

**Service:** Senior Member IEEE and Associate Editor, IEEE Signal Processing Letters

# Overview

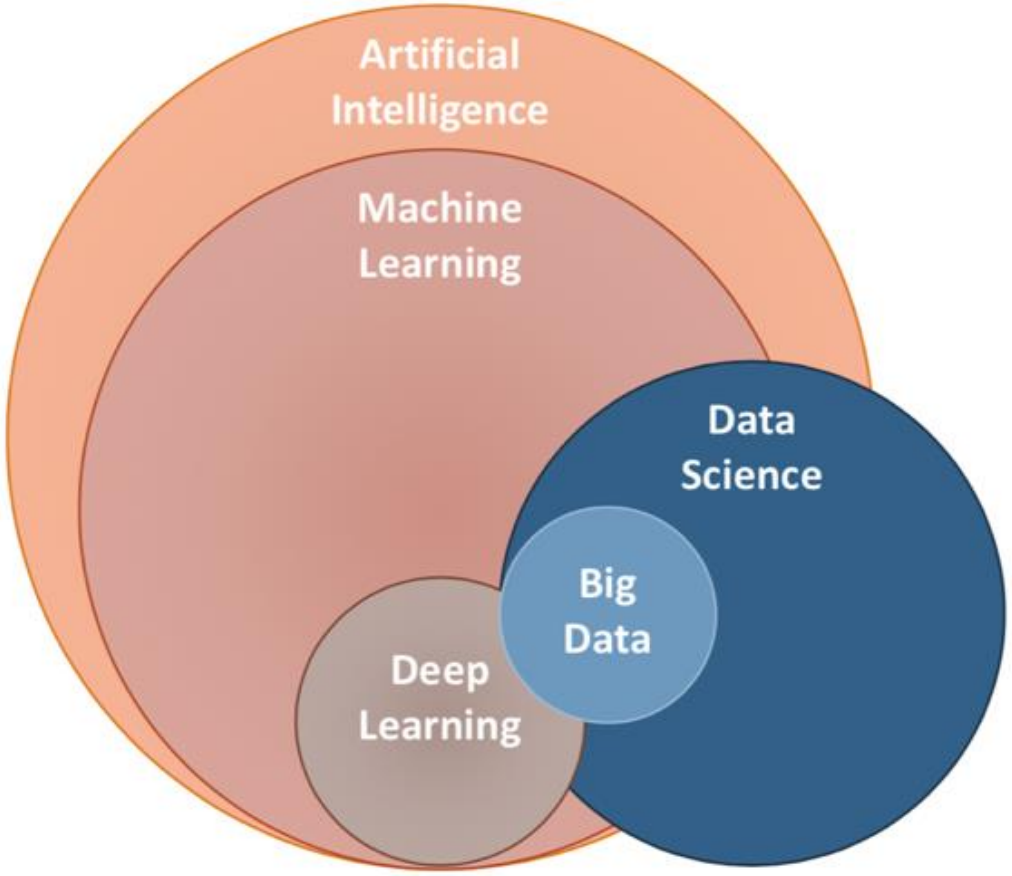


Is it possible to learn artificial intelligence without mathematical knowledge?

# AI and Machine Learning

- Data
  - visualizations, operations, manipulations
- Analysis
  - insights, statistics, inference
- Model
  - representation of process, assumptions
- Learning
  - Optimize model with respect to data

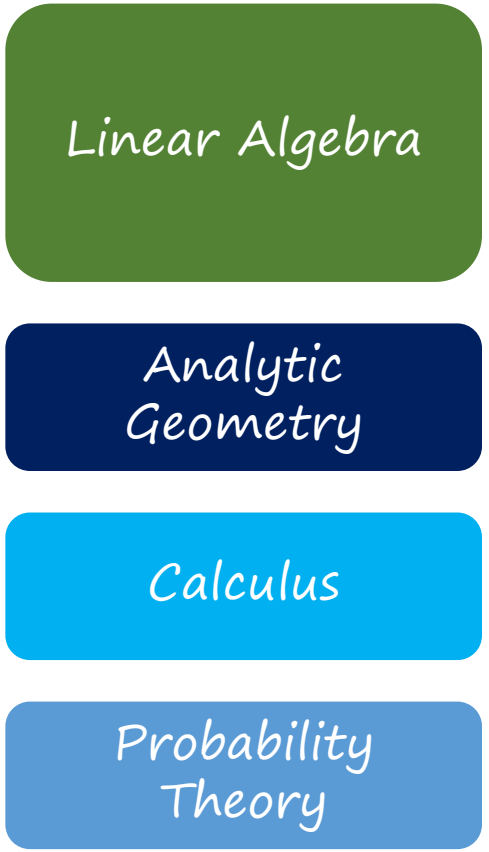

# Course Objectives

- **Foundational Knowledge:** Equip students with a robust understanding of linear algebra, calculus, probability theory, and optimization techniques, which are essential for ML and AI.
- **Practical Applications:** Demonstrate the application of mathematical concepts in real-world ML and AI scenarios, including data processing, model training, and performance evaluation.
- **Analytical Thinking:** Foster analytical thinking by encouraging students to critically assess different algorithms and their suitability for various tasks, considering both theoretical foundations and practical applications.

# Is this course a right choice for you?

**Graduate students**

- Interesting in pursuing a **professional career** focused on the development of AI solutions
- Interested in pursuing AI, Deep Learning (DL) and Machine Learning (ML) in their **grad school (PhD)**
- Want to do fundamental research in the area of AI
- Wish to apply AI in their research work

**Pre-requisites**

- Encouraged to revise Linear Algebra and Probability concepts

# Learning Interface

## Communication and Interaction:

- Class:
  - Saturday, 9:30 to 12:30 Slot
  - Break: 10:45 to 11:10
  - Quiz at 11:10 am
  - Engagement in class

Course Page: https://www.zubairkhalid.org/ai501_2024.html

Slack: Course-related questions or discussions. We will try to respond to the queries ASAP.
https://join.slack.com/t/ai501-mathematicsofai/shared_invite/zt-2pi6n047a-~hn8OE05qqyEnkDyO1jAXA

Office Hours: Posted on course page; distributed throughout the week. TAs during office hours will be available online as well.

# Grading Distribution and Logistics

- 5 Homeworks: 25%
- Quizzes: 25%
- Mid-Exam: 20%
- Final Exam: 30%

# Course Polices

- **Homework Late Policy**
  - 10% per day for 3 days. No submission after 3 days (72 hours)
- **Missed Quiz Policy**
  - No make-up for quiz
- **Plagiarism will be strictly dealt with as per university policies (take it seriously).**
- **Zero Tolerance for Plagiarism and Cheating**
- Re-grading can be requested after grade reporting, within the following time limits:
  - HW and Assignments: 2 days
  - Final Exam: 3 days

# Course Polices

## Harassment Policy

Harassment of any kind is **unacceptable**, whether it be sexual harassment, online harassment, bullying, coercion, stalking, verbal or physical abuse of any kind. Harassment is a very broad term; it includes both direct and indirect behaviour, it may be physical or psychological in nature, it may be perpetrated online or offline, on campus and off campus. It may be one offense, or it may comprise of several incidents which together amount to sexual harassment. It may include overt requests for sexual favours but can also constitute verbal or written communication of a loaded nature. Further details of what may constitute harassment may be found in the LUMS Sexual Harassment Policy, which is available as part of the university code of conduct.

LUMS has a Sexual Harassment Policy and a Sexual Harassment Inquiry Committee (SHIC). Any member of the LUMS community can file a formal or informal complaint with the SHIC. If you are unsure about the process of filing a complaint, wish to discuss your options or have any questions, concerns, or complaints, please write to the Office of Accessibility and Inclusion (OAI, oai@lums.edu.pk) and SHIC (shic@lums.edu.pk) —both of them exist to help and support you and they will do their best to assist you in whatever way they can.

**To file a complaint, please write to harassment@lums.edu.pk.**

# Course Polices

## Help related to equity and Belonging at SSE

SSE's Council on Equity and Belonging is committed to devising ways to provide a safe, inclusive, and respectful learning, living, and working environment for its students, faculty, and staff.
For help related to any such issue, please feel free to write to any member of the school council for help or feedback.

## Mental Health Support at LUMS

-

For matters relating to counselling, kindly email student.counselling@lums.edu.pk, or visit https://osa.lums.edu.pk/content/student-counselling-office for more information.

You are welcome to write to me or speak to me if you find that your mental health is impacting your ability to participate in the course. However, should you choose not to do so, please contact the Counselling Unit and speak to a counsellor or speak to the OSA team and ask them to write to me so that any necessary accommodations can be made.

# Suggested Reference Books

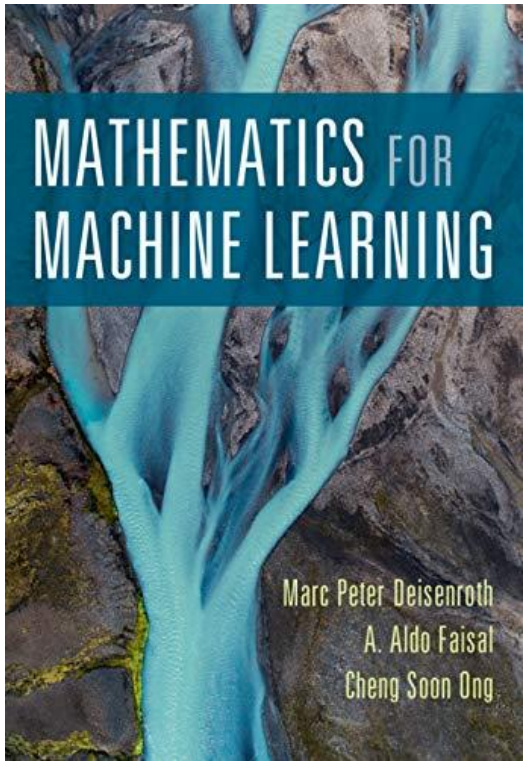


Primary

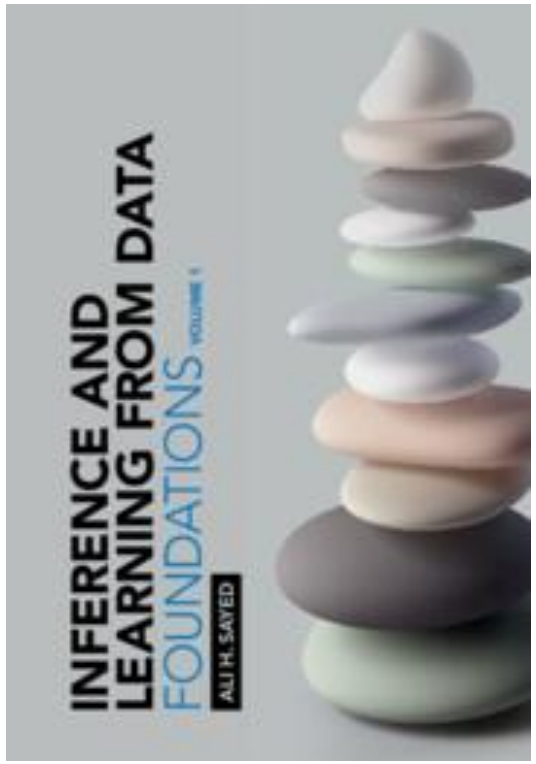


Reference

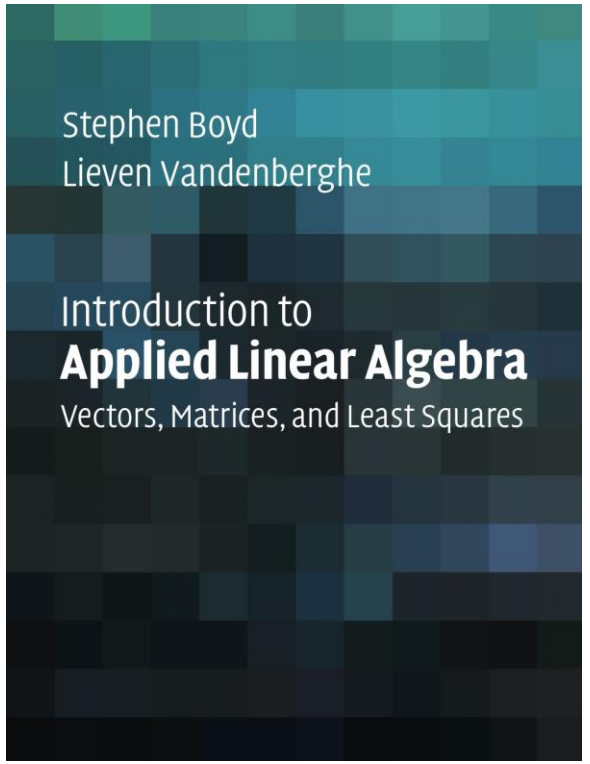


Reference

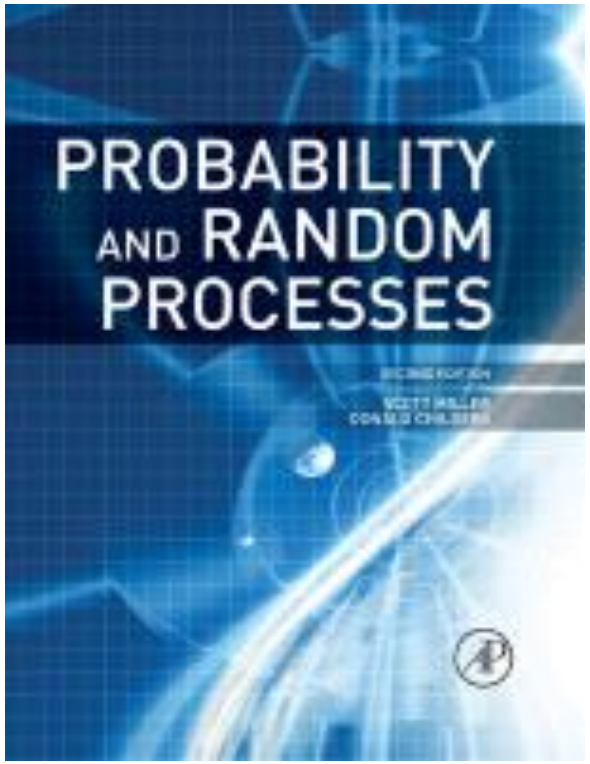


Reference

- Lecture Notes/Slides will be shared.

*"As to methods, there may be a million and then some, but principles are few. The man who grasps principles can successfully select his own methods."*

**Ralph Waldo Emerson**

# Outline

- Course Overview, Logistics
- AI Overview
- Introduction to Machine Learning

# AI Overview

## What is Intelligence?

"Intelligence is the ability to **learn**, **understand**, and **apply** knowledge to adapt to new situations, reason, and solve problems."

- Human Intelligence
  - Problem-solving and reasoning
  - Learning from Experience
  - Adaptability
  - Emotional Intelligence

# AI Overview

## What is Artificial Intelligence (AI)?

"AI refers to the simulation of **human intelligence** in machines that are programmed to think and learn like humans."

- Narrow AI vs General AI
- SuperIntelligence

The term AI was coined in 1956 by John McCarthy at a conference.

# AI Overview

## History of AI – Turing Test

Alan Turing (1950) "Computing machinery and intelligence"

- **Can machines think or behave intelligently?**
- The Turing Test (imitation game) was designed as a way to judge the success or otherwise of an attempt to produce a thinking computer.
- Predicted that by 2000, a machine would have a 30% chance of fooling a person for 5 minutes
- What do you think?

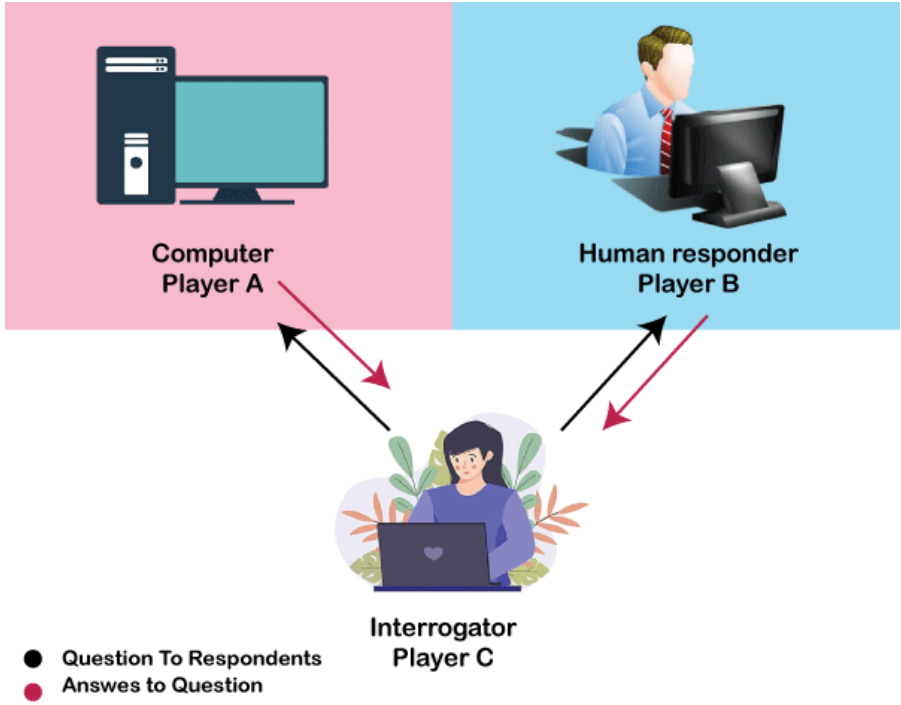

# AI Overview

## History of AI – 1960s to 1990s

- Since the 1950s, much of the initial optimism surrounding AI has given way to a more realistic perspective.

- Key areas of focus include:
  - Machine learning
  - Multi-agent systems
  - Computer vision
  - Planning
  - Game Theory
  - Optimization, …

- 1997 (IBM's Deep Blue), 2012 (Deep Learning breakthrough).

# Applications of AI

## AI in Healthcare

- **AI-powered Diagnostics**

- AI systems can analyze medical images, such as X-rays, MRIs, and CT scans, to detect diseases like cancer, heart conditions, and more.

- These systems can identify patterns that might be missed by human eyes, resulting in earlier and more accurate diagnoses.

- **Example:** Google's DeepMind developed an AI system capable of diagnosing eye diseases with an accuracy that matches world-leading doctors.

# Applications of AI

## AI in Healthcare

- **Predictive Analytics**

- AI can analyze patient data and predict potential health risks, such as the likelihood of developing chronic diseases like diabetes or heart disease. This helps doctors create personalized treatment plans and intervene early.

- **Example:** IBM Watson Health analyzes vast datasets to predict disease progression and recommend personalized treatment options.

# Applications of AI

## AI in Healthcare

- **Robotic Surgery**

- AI-powered robotic systems assist in precision surgeries, allowing for less invasive procedures, faster recovery times, and reduced human error.

- **Example:** The da Vinci Surgical System uses AI to enhance the precision and control of surgeons during minimally invasive procedures.

# Applications of AI

## AI in Healthcare

- **Drug Discovery and Development**

- AI accelerates the drug discovery process by analyzing complex biological data to identify new drug candidates, significantly reducing the time and cost involved in developing new treatments.

- **Example:** AI was used by the pharmaceutical company Insilico Medicine to identify potential new compounds for treating fibrosis, which traditionally would have taken years.

# Applications of AI

## AI in Healthcare

- **Virtual Health Assistants**

- AI chatbots and virtual assistants provide personalized health advice, answer questions, and help patients manage chronic conditions by monitoring symptoms and offering medication reminders.

- **Example:** AI-driven platform allows patients to receive medical advice through their smartphones based on the symptoms they input.

# Applications of AI

## AI in Finance

- Fraud Detection
- Algorithmic Trading
- Risk Management
- Personalized Banking
- Credit Scoring and Loan Decisions
- Financial Forecasting

# Applications of AI

## AI in Education

- Personalized Learning Systems
- Virtual Tutors and Chatbots
- AI in Assessments and Grading
- AI in Special Education

# Applications of AI

## AI in Autonomous Vehicles

- Deep Learning for Object Detection
- Decision Making and Planning (Reinforcement Learning)
- Simultaneous Localization and Mapping (SLAM)
- Path Planning and Navigation
- To name a few…

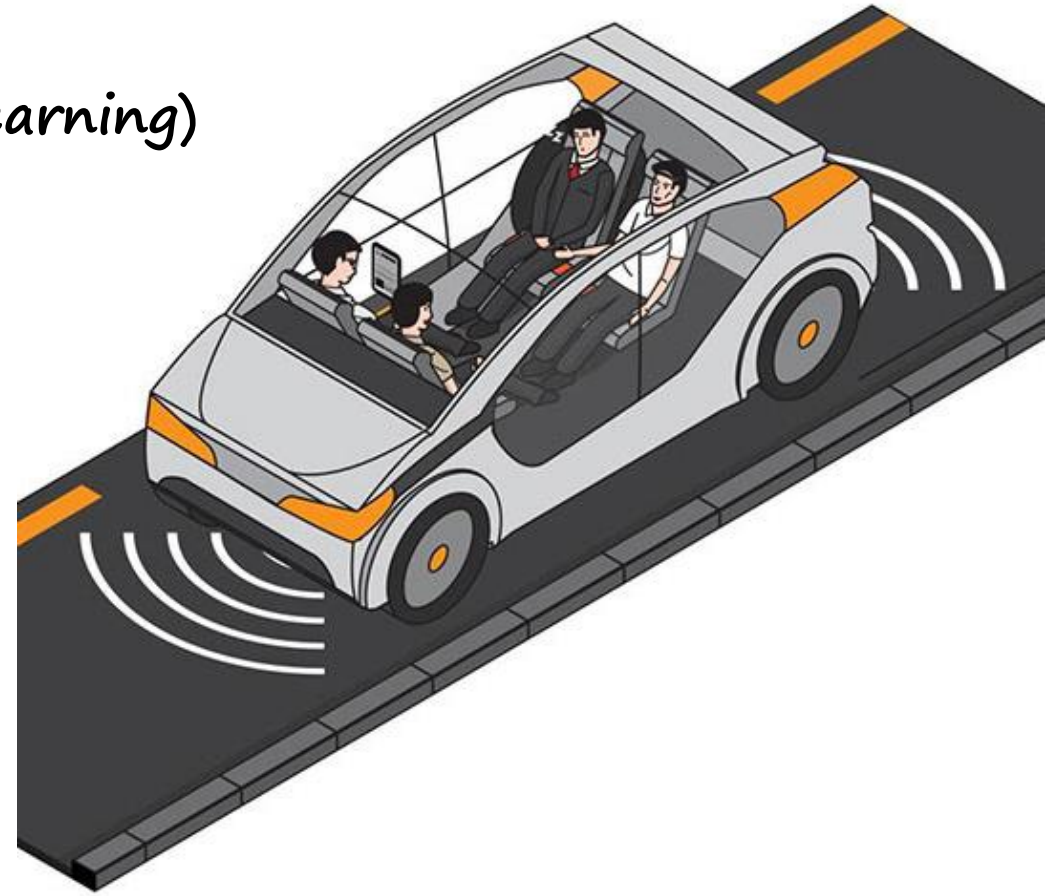

# Applications of AI

## AI in Retail

- Personalized product recommendations
- AI-powered chatbots and virtual assistants
- Inventory management and demand forecasting
- Dynamic pricing
- Visual search and image recognition
- Customer sentiment analysis

# Applications of AI

## AI in Agriculture

- Precision farming
- Crop monitoring and disease detection
- Soil health monitoring and analysis
- Automated machinery and robotics
- Predictive analytics and yield forecasting

# Applications of AI

## AI in NLP

- Virtual assistants (e.g., Siri, Alexa)
- Language translation
- Text generation and summarization
- Speech recognition and transcription
- Sentiment analysis
- Chatbots for customer support

# Applications of AI

## AI in Manufacturing

- Predictive maintenance
- Quality control and defect detection
- Supply chain optimization
- Autonomous robots and cobots (collaborative robots)
- Inventory management and optimization
- Production scheduling and process optimization

# Applications of AI

## AI in Urban Development

- Traffic management and optimization
- Smart energy grids and resource management
- Road safety and surveillance
- Waste management optimization
- Smart infrastructure and building management
- Environmental monitoring

# Centre for Urban Informatics, Technology and Policy (CITY)

## *Saving our cities through technology and data-driven policy*

| VISION STATEMENT | *To become an internationally recognized regional hub of innovation in urban informatics, technology and data-driven policy-making for sustainable urban development.* |
|---|---|

# Centre for Urban Informatics, Technology and Policy (CITY)

## Growth

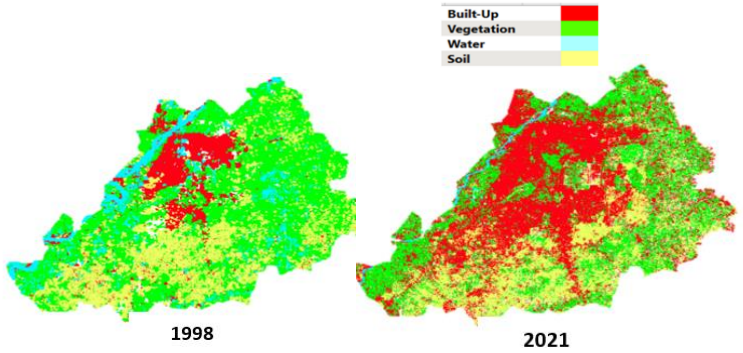


**Lahore's Built-up Area. 4X (1998 – 2021)**

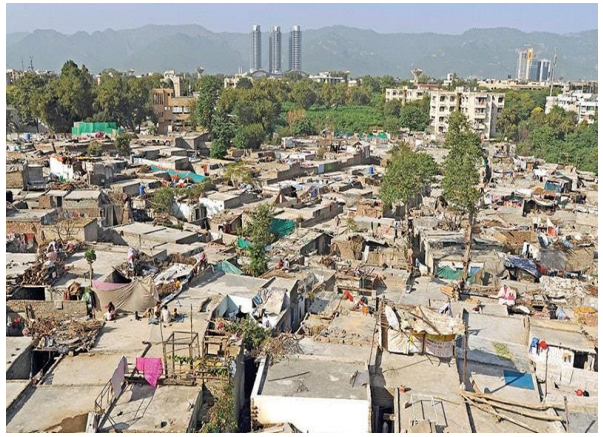

*Murtaza Haider, "Pakistan's urban policy turning cities into slums", Dawn, April 2015.*

## Mobility

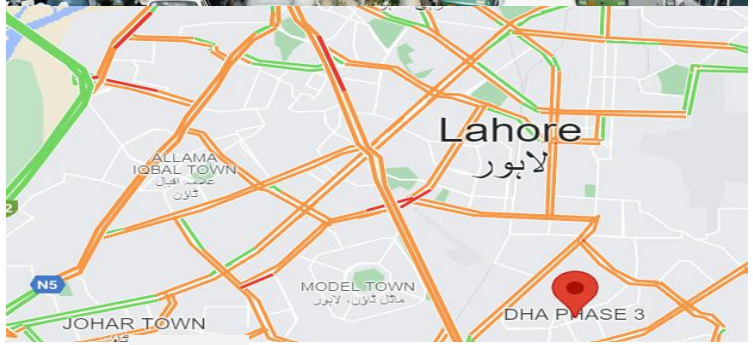


3x (2010-2020)

2x (2010-2020)

## Environment and Health

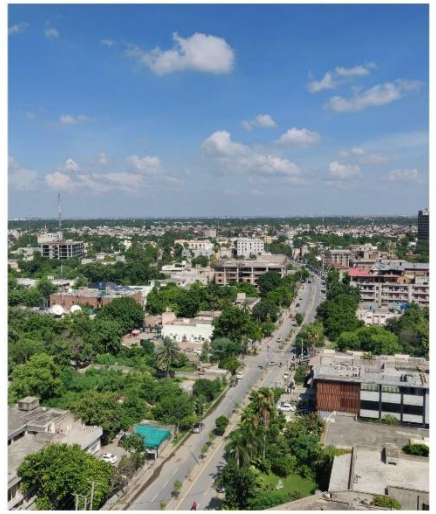

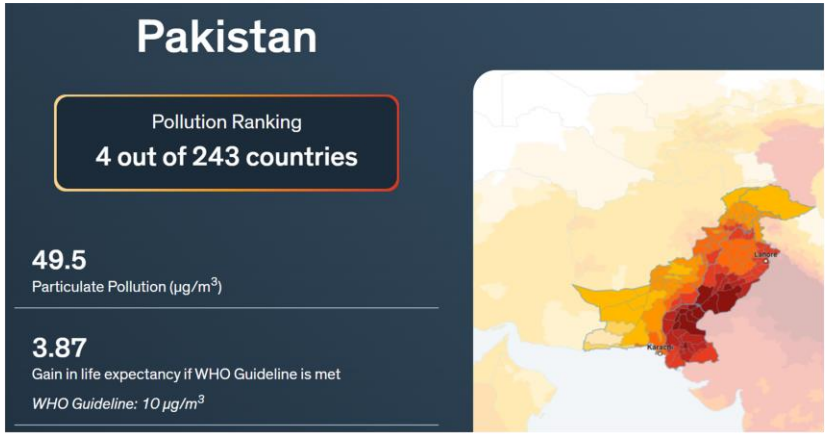


LUMS
A Not-for-Profit University

# Urban Mobility: Road Network

Data Layers

Google Earth

Google Maps Platform

Open Street Maps

Manual Labeling

Machine Learning and Image Processing

Road Network Lahore

Objective:
Creating a dense road network with associated attributes like speed, distance, lanes, etc.

Road Network

# Urban Mobility: Road Network

Analysis Tools:

QGIS, Python, OpenCV, and SciPy

Road Network

# Urban Mobility: Road Network

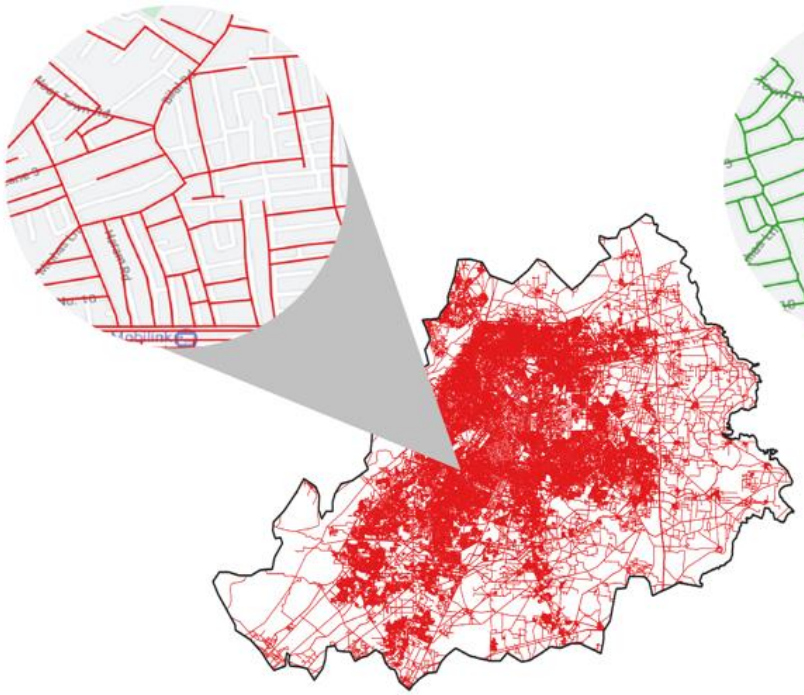

OSM Road Network

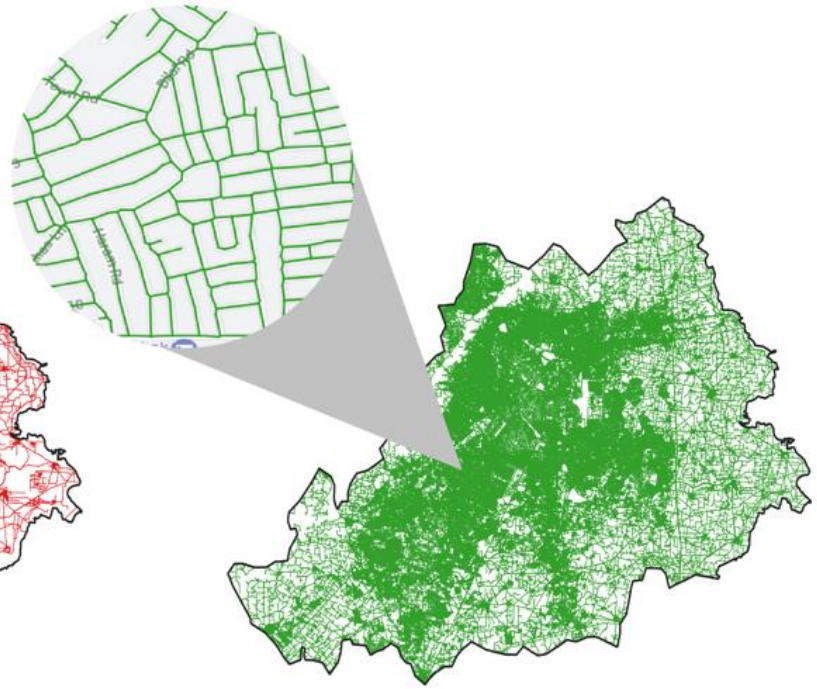

Our Road Network

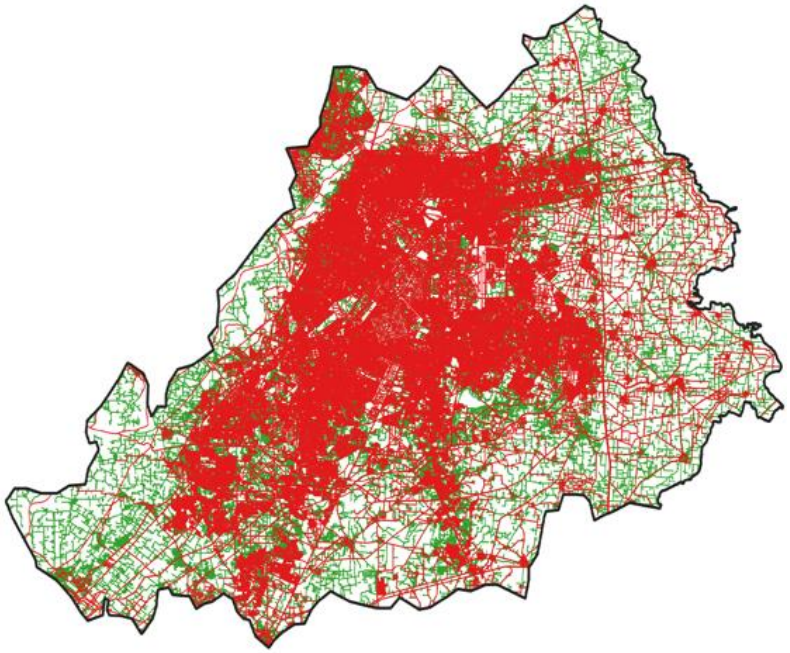

Overlaid

# Urban Mobility: Traffic Informatics

Road monitoring using camera

Database/Storage

AI model

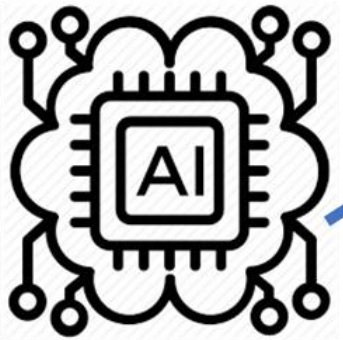


Dashboard

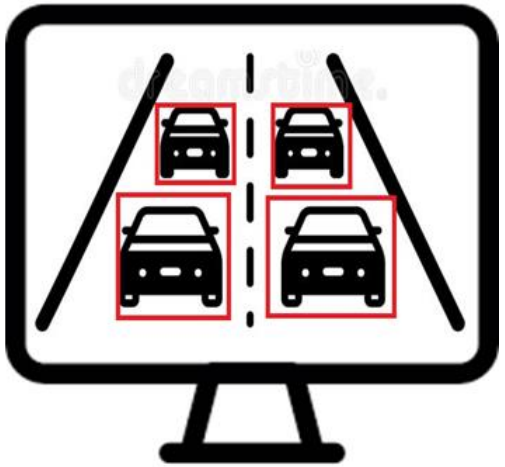

AI based system for real time traffic flow density

Traffic Flows Density

# Urban Mobility: Traffic Informatics


Prototype Sample

Traffic Flows Density

LUMS
A Not-for-Profit University

# Urban Sprawl: Land Surface Temperature

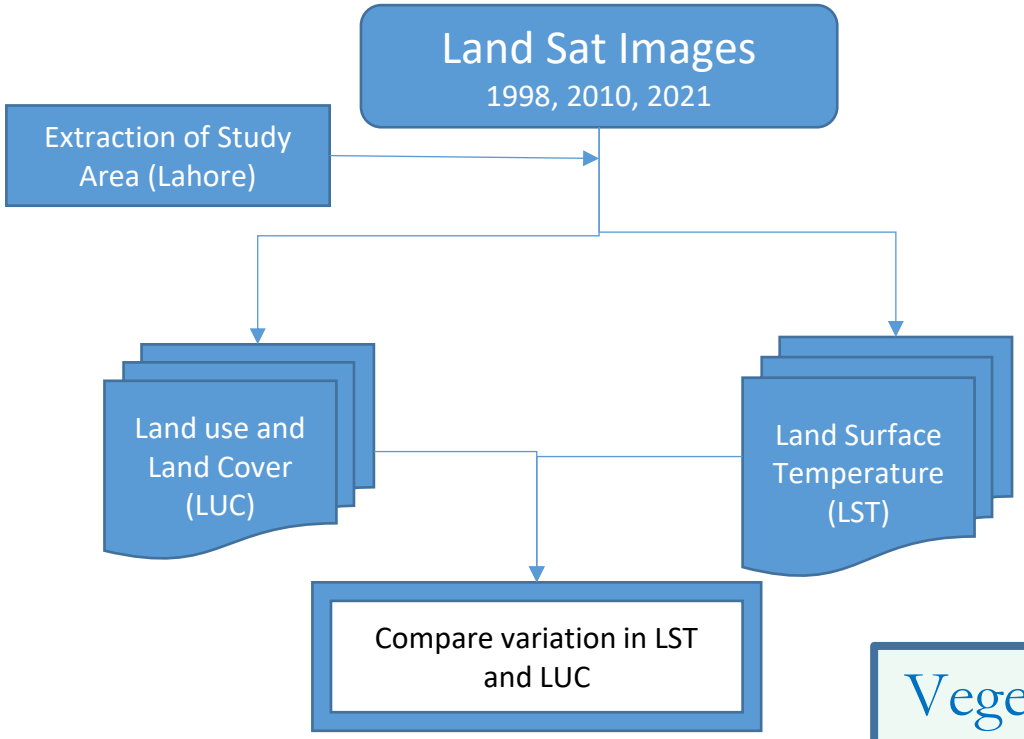


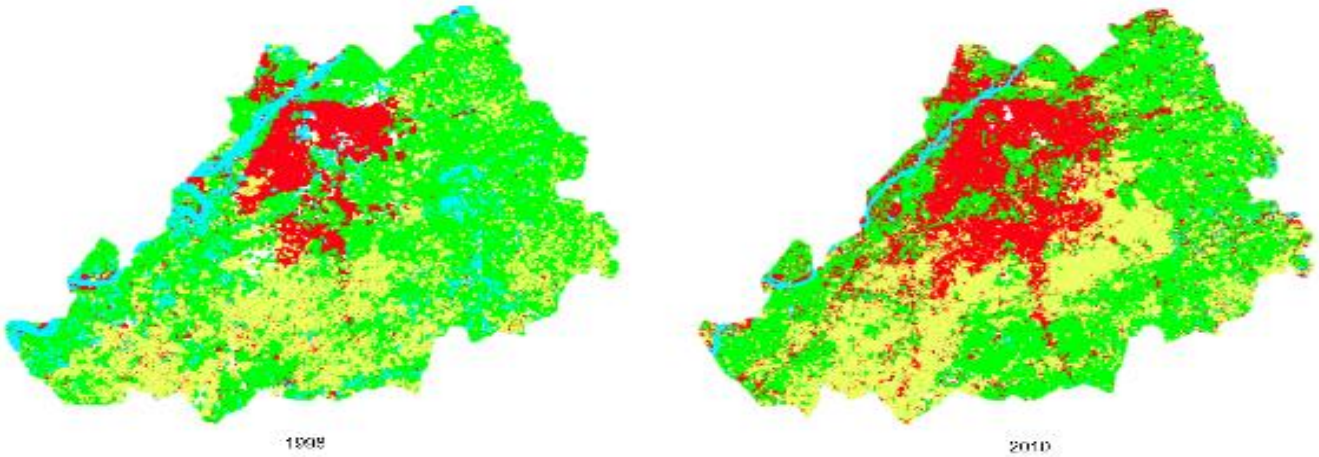


Vegetation and waterbodies declined; Built-up use grew significantly

An average increase of 0.47 Celsius per year between 1998 and 2021

LUMS
A Not-for-Profit University

# Data-driven Urban Forestation: Tree Cover Estimation

Problem: **Estimate green spaces and tree cover in urban areas**

**Why?**

Green spaces – significant contributors to ecosystem services: air purification and carbon sequestration

**Challenges**

- Bottleneck: absence of spatially explicit data
- Huge variations in satellite imagery

**Existing Methodologies**

- Current methodologies focus on forest mapping
- A gap exists in urban tree cover estimation in developing countries with low resolution imagery

# Data-driven Urban Forestation: Tree Cover Estimation

Tree Cover Estimation using Machine Learning and Feature Selection on Sentinel-2 Data

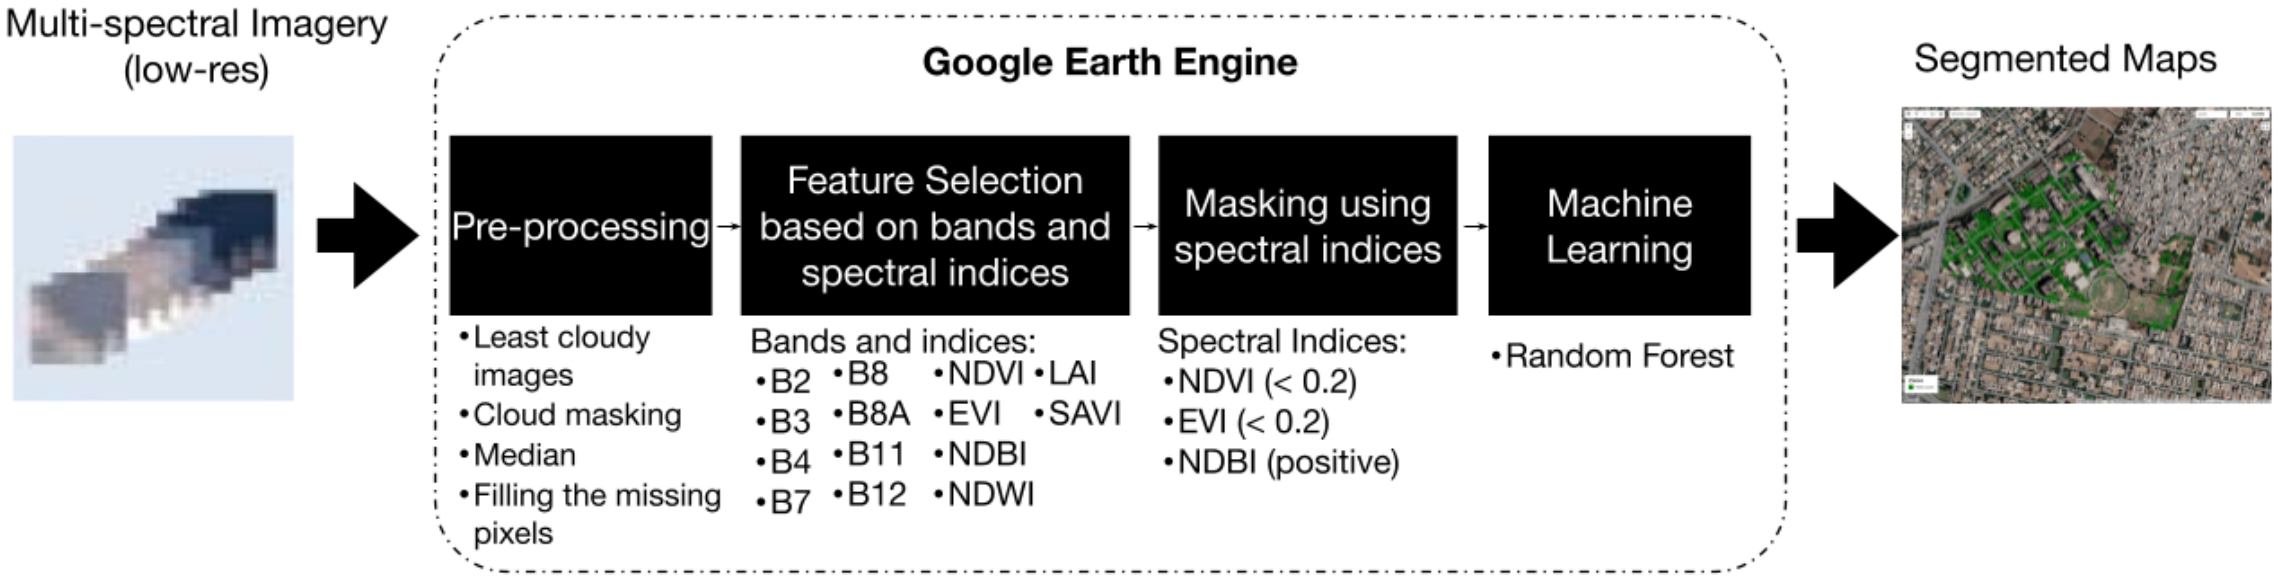


**U. Nazir, M. Uppal, M. Tahir, and Z. Khalid, "Feature Selection on Sentinel-2 Multi-Spectral Imagery for Efficient Tree Cover Estimation," 2023 IEEE International Geoscience and Remote Sensing Symposium (IGARSS 2023), Pasadena, California**

# Data-driven Urban Forestation: Tree Cover Estimation

LUMS Ground Truth

Proposed

ESA WorldCover

| Model | Pred. area (acres) | Masking | Spectral indices | Pixel-wise Test Accuracy (%) | Kappa Score |
|---|---|---|---|---|---|
| RF-spectral-bands | 29.5 | No | No | 0.93 | 0.81 |
| RF-spectral-indices | 28 | No | Yes | 0.95 | 0.88 |
| **Proposed** | 25 | Yes | Yes | **0.99** | **0.92** |
| ESA WorldCover Product [5] | 16 | - | - | 0.74 | - |
| DeepLabv3 [6] | 28 | No | No | 0.80 | - |

# Outline

- Course Overview, Logistics
- AI Overview
- Introduction to Machine Learning

# Machine Learning Overview

## What is Machine Learning?

- Automating the process of automation
- Getting computers to program themselves


**Traditional Programming**

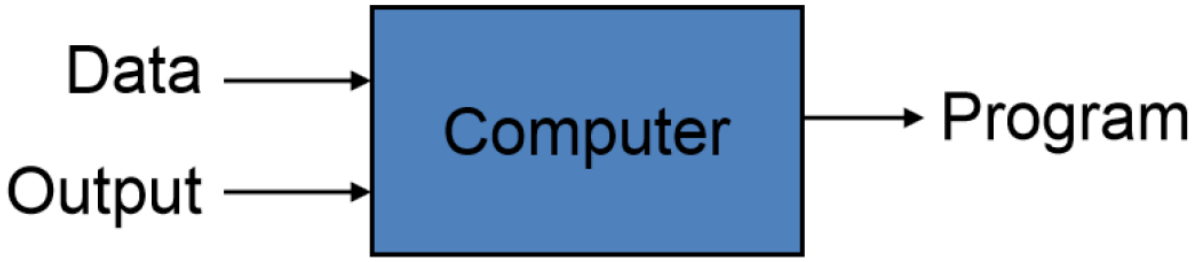


**Machine Learning**

Given examples (training data), make a machine learn system behavior or discover patterns

# Machine Learning Overview

*Classical Example: Recognize hand-written 2!*

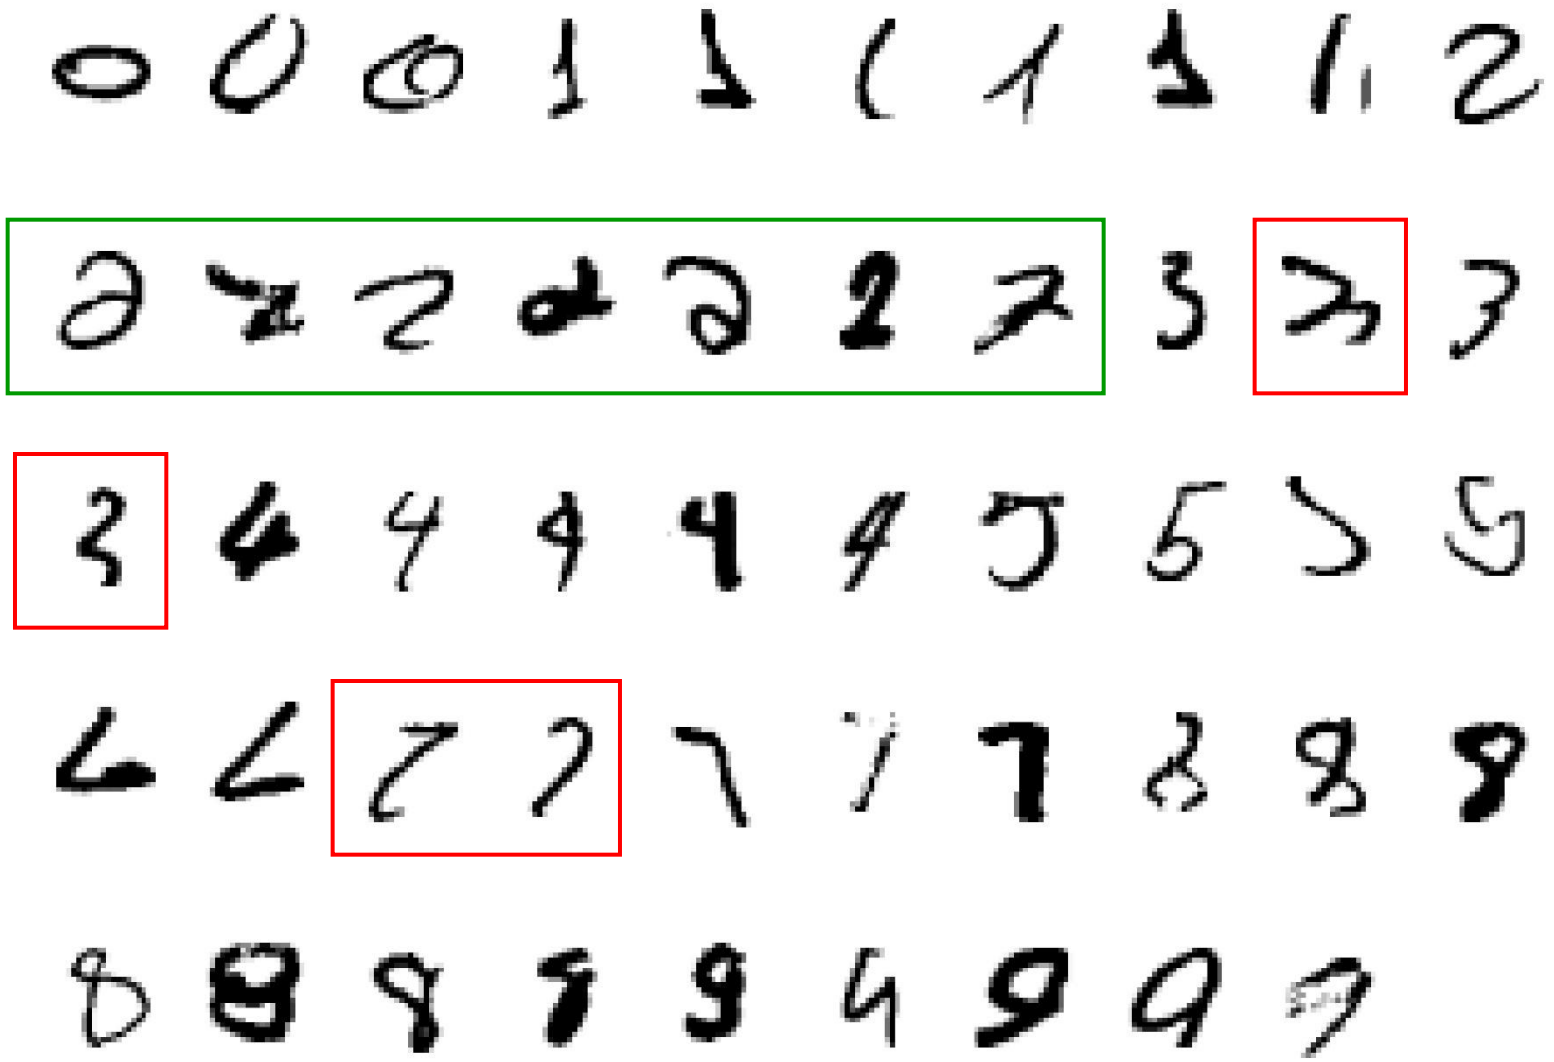

# Machine Learning: Overview

## What is Machine Learning?

*Given examples (training data), make a machine learn system behavior or discover patterns*

**Data**

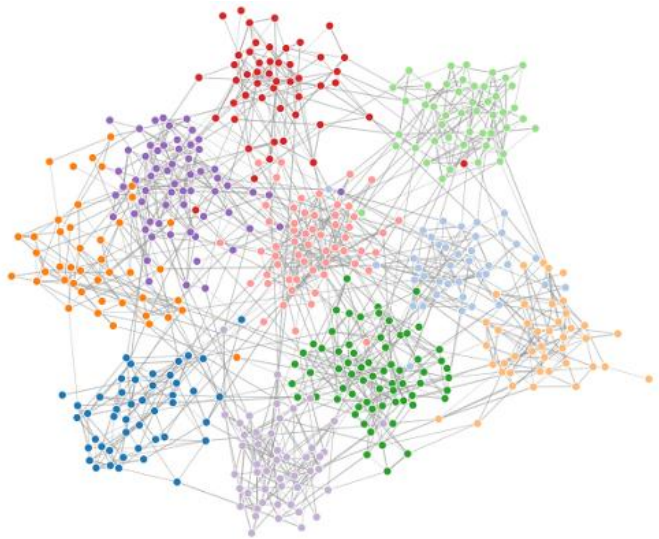

**Algorithm**

**Model**

$$f(\mathbf{x})$$

**Given to us**

**We need to design it**

**Final output which enables us to make predictions**

# Machine Learning: Overview

## Algorithms vs Model

- Linear regression algorithm produces a model, that is, a vector of values of the coefficients of the model.

- Decision tree algorithm produces a model comprised of a tree of if-then statements with specific values.

- Neural network along with backpropagation + gradient descent: produces a model comprised of a trained (weights assigned) neural network.

# Machine Learning: Overview

## *Nature of ML Problems*

### 1. Supervised Learning

The learning algorithm would receive a set of inputs along with the corresponding correct outputs to train a model

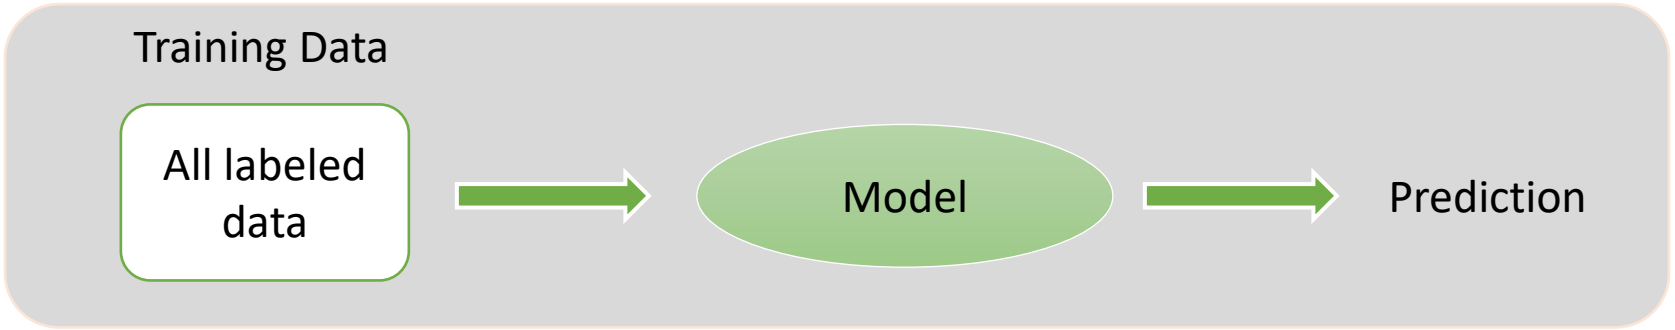

# Supervised Learning

## *Regression*

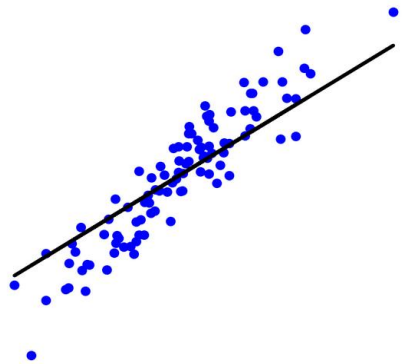

**Regression:** Quantitative Prediction on a continuous scale

**Examples: Prediction of**

- Age of a person from his/her photo
- Price of 10 Marla, 5-bedroom house in 2050
- USD/PKR exchange rate after one year
- Efficacy of vaccine or medicine
- Average temperature/Rainfall during monsoon
- Cumulative score in AI501 course
- Probability of decrease in the electricity prices in Pakistan
- No. of steps per day

What do all these problems have in common?

Continuous outputs

Predicting continuous outputs is called regression

# Supervised Learning

## *Classification*

**Classification:** Given a data sample, predict its class **(discrete)**

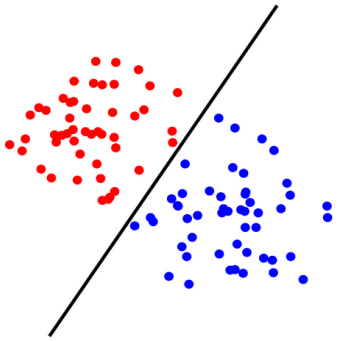

**Examples: Prediction of**

- Gender of a person using his/her photo or hand-writing style
- Spam filtering
- Object or face detection in a photo
- Temperature/Rainfall normal or abnormal during monsoon
- Letter grade in AI501 course
- Decrease expected in electricity prices in Pakistan next year
- More than 10000 Steps taken today

What do all these problems have in common?

Discrete outputs: Categorical

Yes/No (Binary Classification)

Multi-class classification: multiple classes

Predicting a categorical output is called classification

# Machine Learning: Overview

## *Nature of ML Problems*

### 2. Unsupervised Learning

*The learning algorithm would receive unlabeled raw data to train a model and to find patterns in the data*

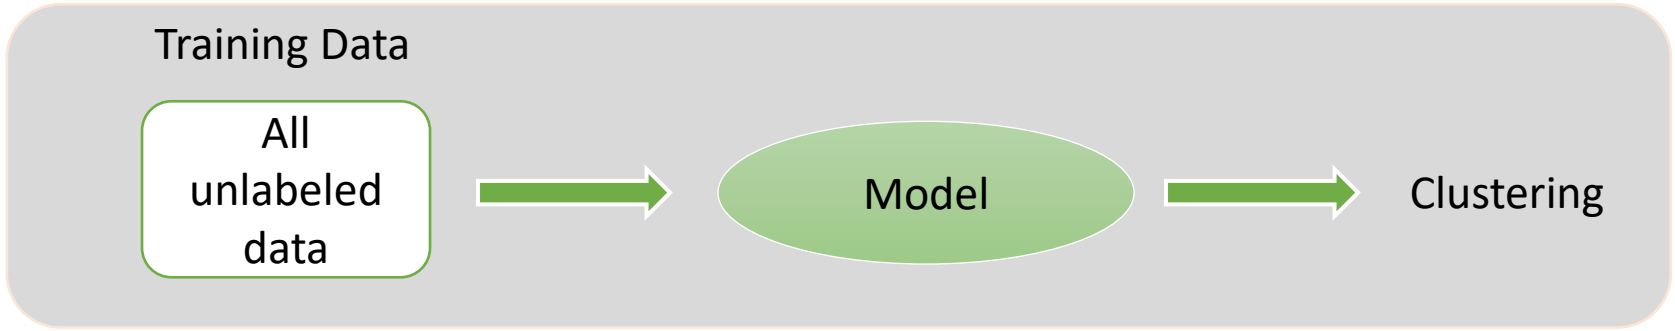


# Machine Learning: Overview

## *Nature of ML Problems*

### 3. Semi-supervised Learning

- *The learning algorithm receives labeled and unlabeled raw data to train a model*
- *Main objective is to efficiently accommodate the unlabeled data*

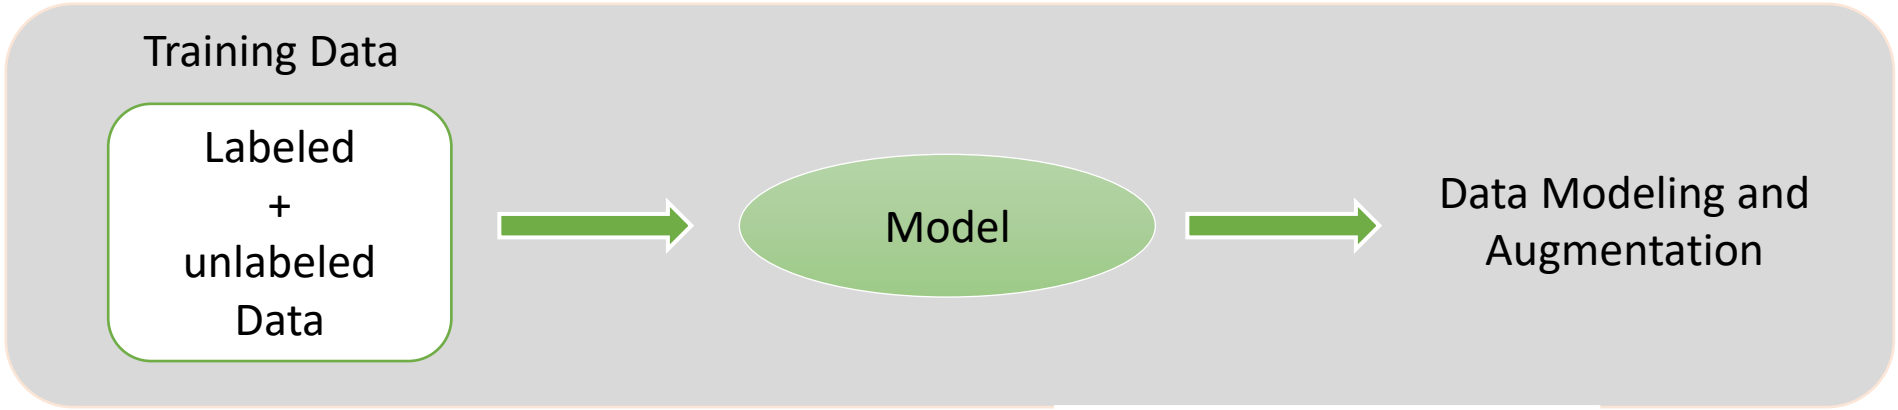


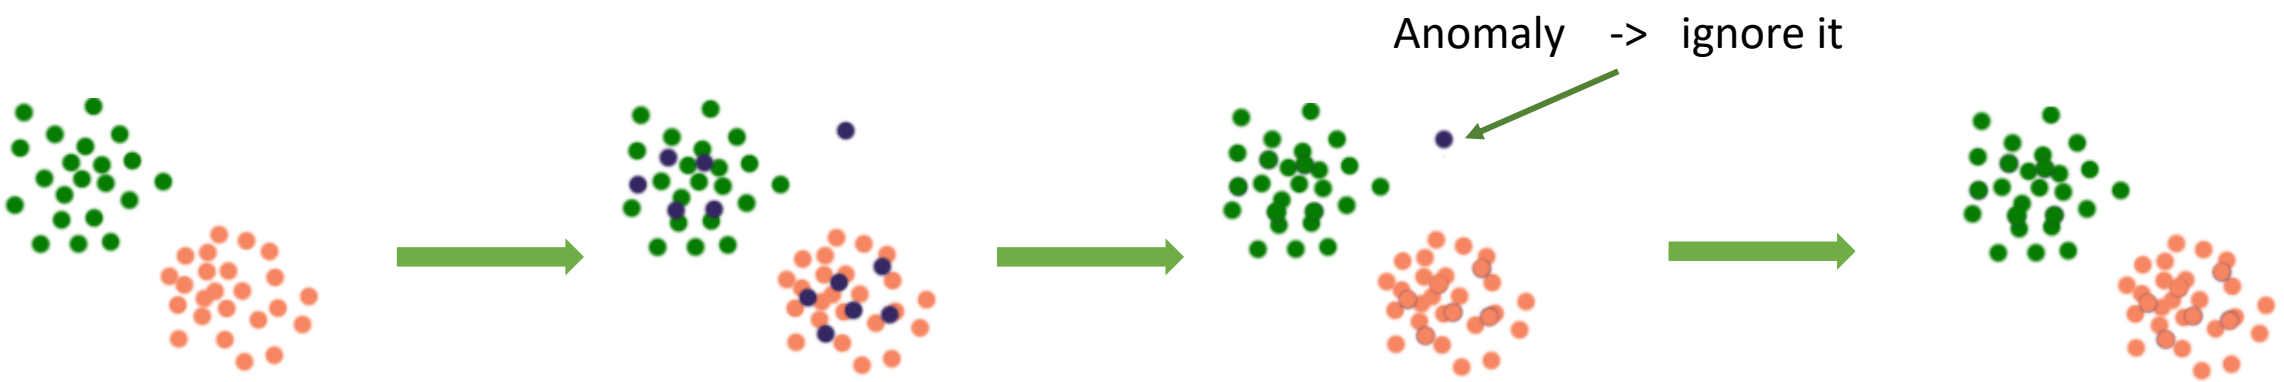


Use labeled data to build a model

Place unlabeled data with model

Use the model to label the unlabeled data

Fit the model again with the combined data

# Machine Learning: Overview

## *Training Data Collection*

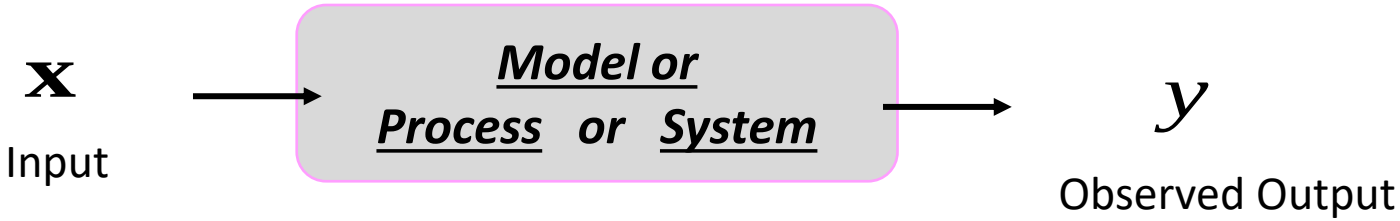


***PROCESS*** or ***SYSTEM*** : Underlying physical or logical phenomenon which maps our input data to our observed output

Collect the training data by observing our unknown **PROCESS** or **SYSTEM**

# Machine Learning: Overview

## *Example Systems*

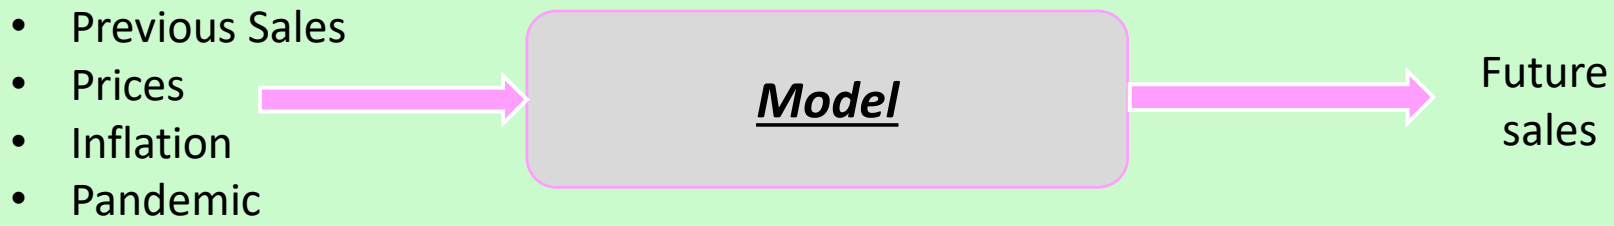


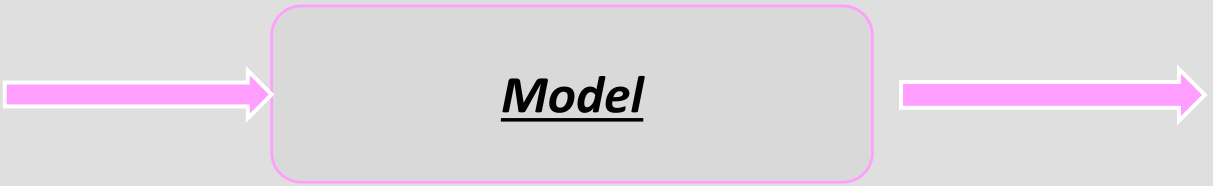

# Machine Learning: Overview

## *Example Systems*

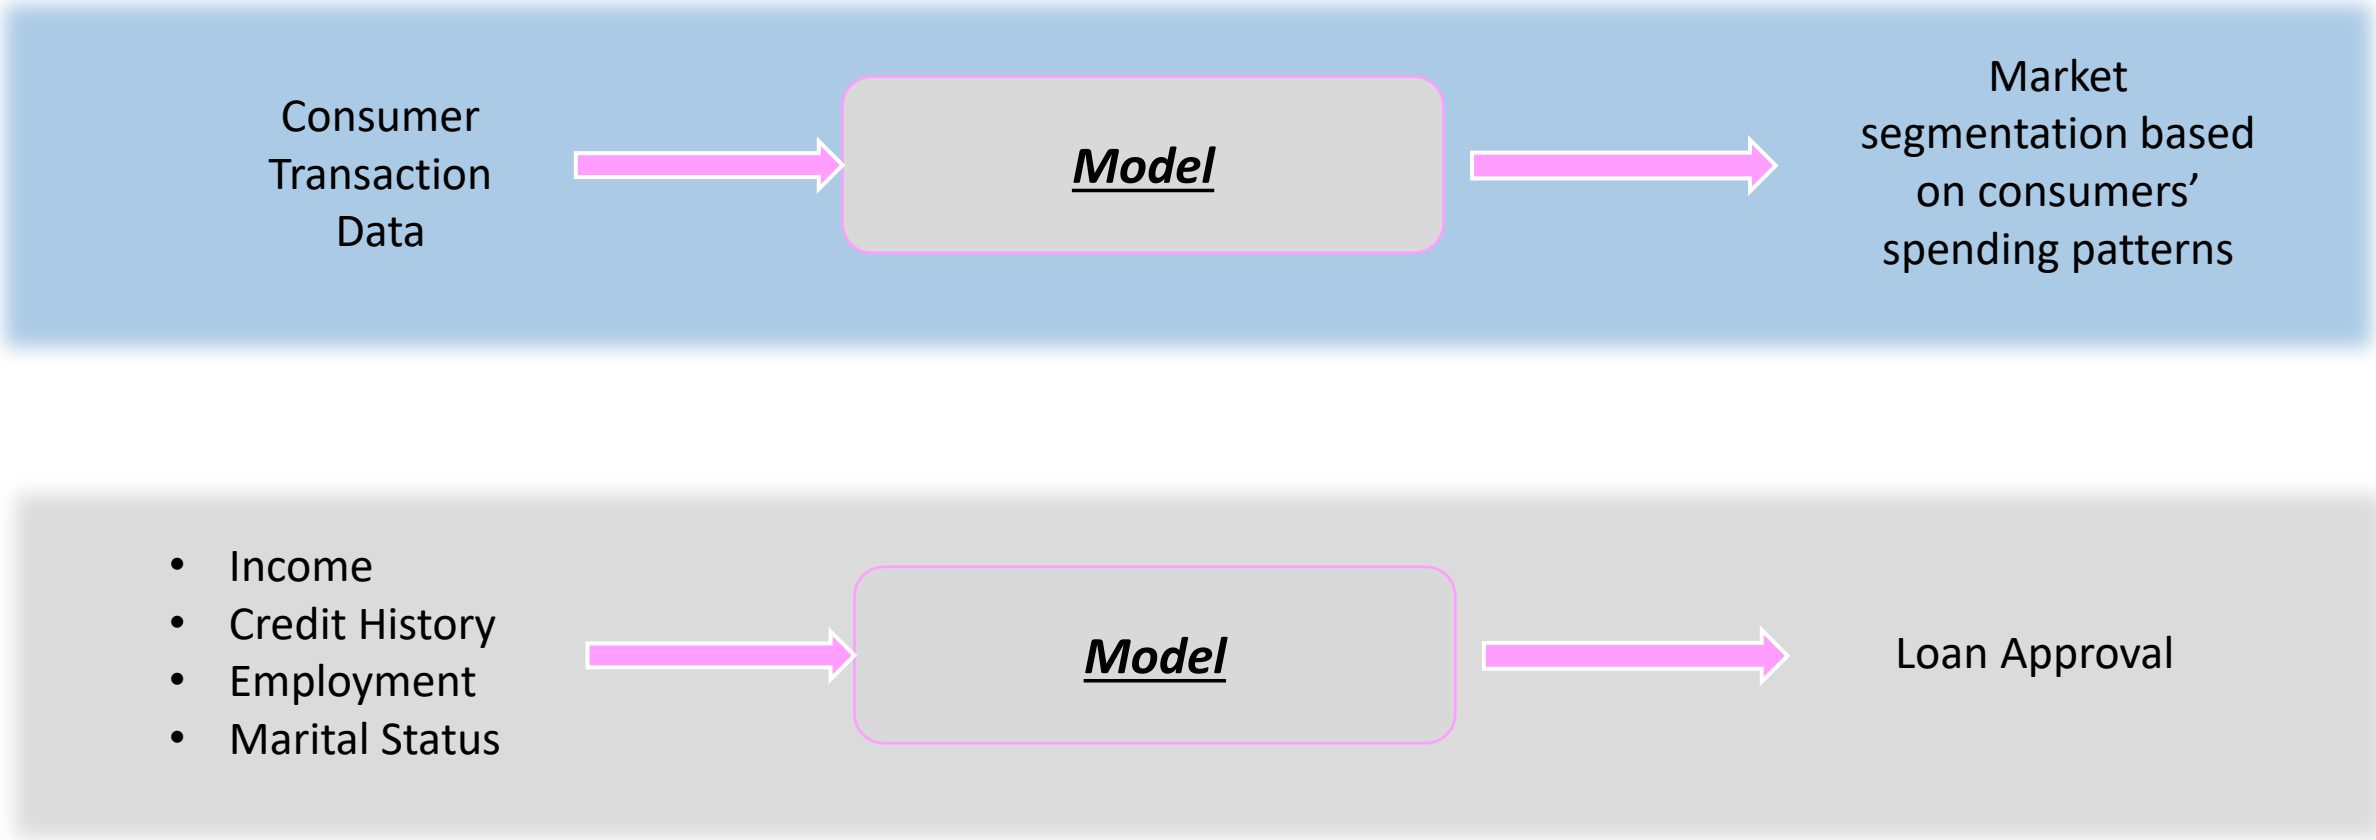

# Machine Learning: Overview

*Typical Flow*

Problem Nature Identification

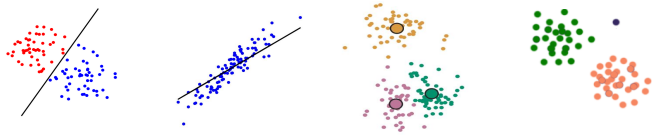

Training Data Collection

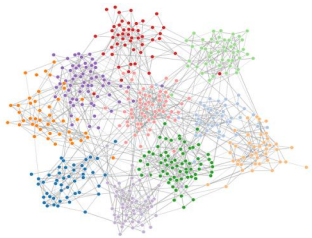

Model Training

$$f(\mathbf{x})$$

Model Validation

Model Deployment

# Supervised Learning Setup

## Nomenclature

In these regression or classification problems, we have

- Inputs – referred to as Features
- Output – referred to as Label
- Training data – (input, output) for which the output is known and is used for training a model by ML algorithm
- A Loss, an objective or a cost function – determines how well a trained model approximates the training data
- Test data – (input, output) for which the output is known and is used for the evaluation of the performance of the trained model

# Supervised Learning Setup

## Nomenclature - Example

Predict Stock Index Price

- *Features (Input)*
- *Labels (Output)*
- *Training data*

| Interest_Rate | Unemployment_Rate | Stock_Index_Price |
|---|---|---|
| 2.75 | 5.3 | 1464 |
| 2.5 | 5.3 | 1394 |
| 2.5 | 5.3 | 1357 |
| 2.5 | 5.3 | 1293 |
| 2.5 | 5.4 | 1256 |
| 2.5 | 5.6 | 1254 |
| 2.5 | 5.5 | 1234 |
| 2.25 | 5.5 | 1195 |
| 2.25 | 5.5 | 1159 |
| 2.25 | 5.6 | 1167 |
| 2 | 5.7 | 1130 |
| 2 | 5.9 | 1075 |
| 2 | 6 | 1047 |
| 1.75 | 5.9 | 965 |
| 1.75 | 5.8 | 943 |
| 1.75 | 6.1 | 958 |
| 1.75 | 6.2 | 971 |
| 1.75 | 6.1 | 949 |
| 1.75 | 6.1 | 884 |
| 1.75 | 6.1 | 866 |
| 1.75 | 5.9 | 876 |
| 1.75 | 6.2 | ? |
| 1.75 | 6.2 | ? |
| 1.75 | 6.1 | ? |

# Supervised Learning Setup

## Example

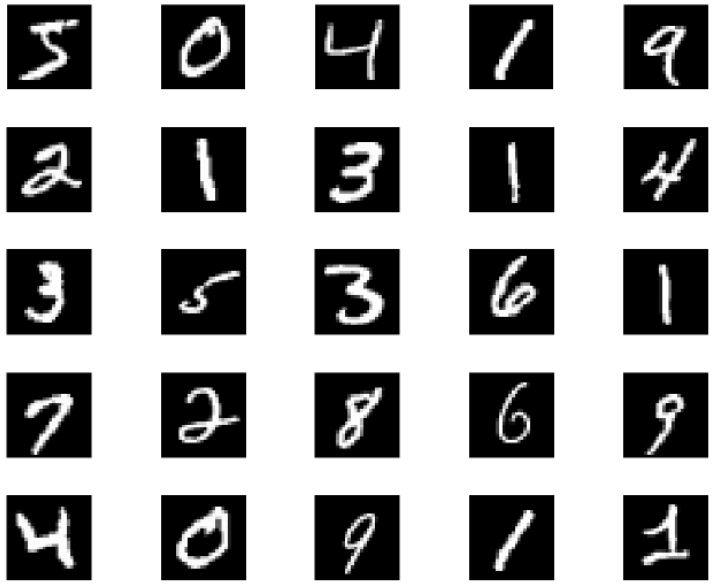

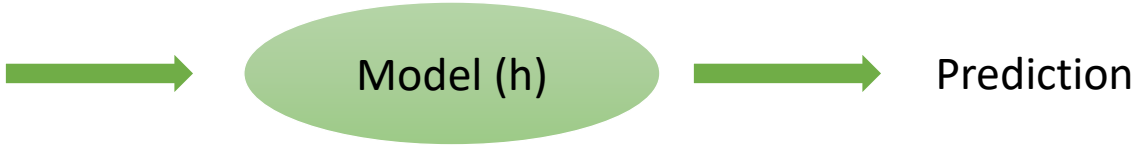


**MNIST Data:**

- Each sample 28x28 pixel image
- 60,000 training data
- 10,000 testing data

# How are you feeling?

**A**


**B**


**C**


**D**


Scan QR code to respond

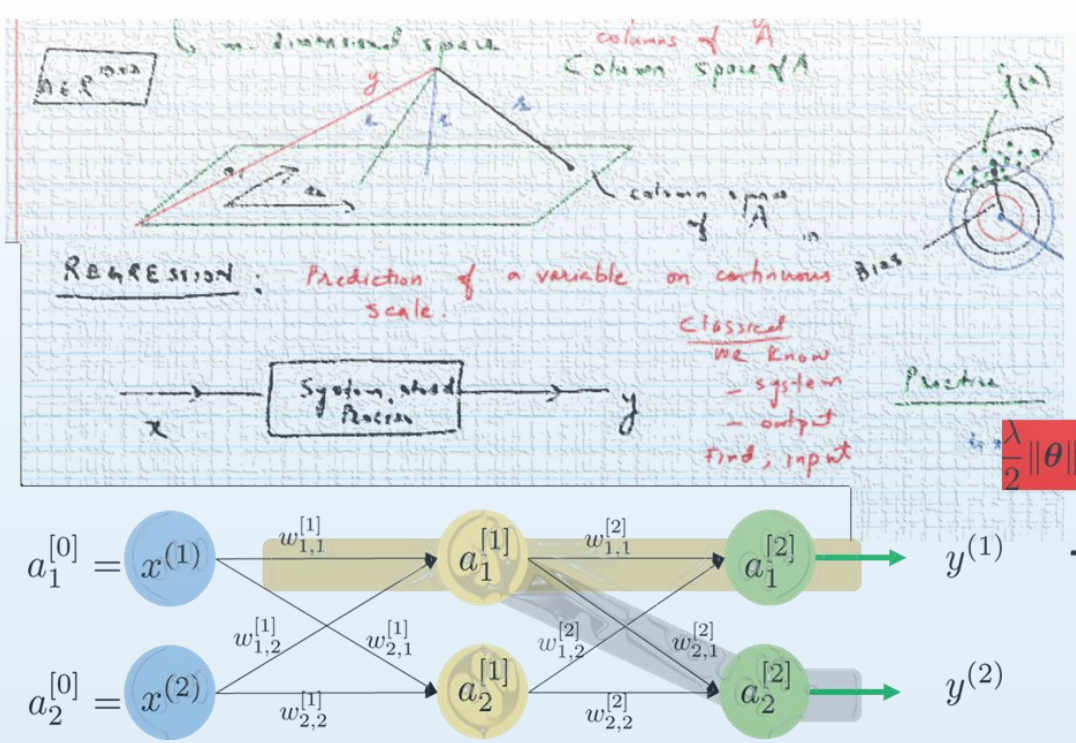


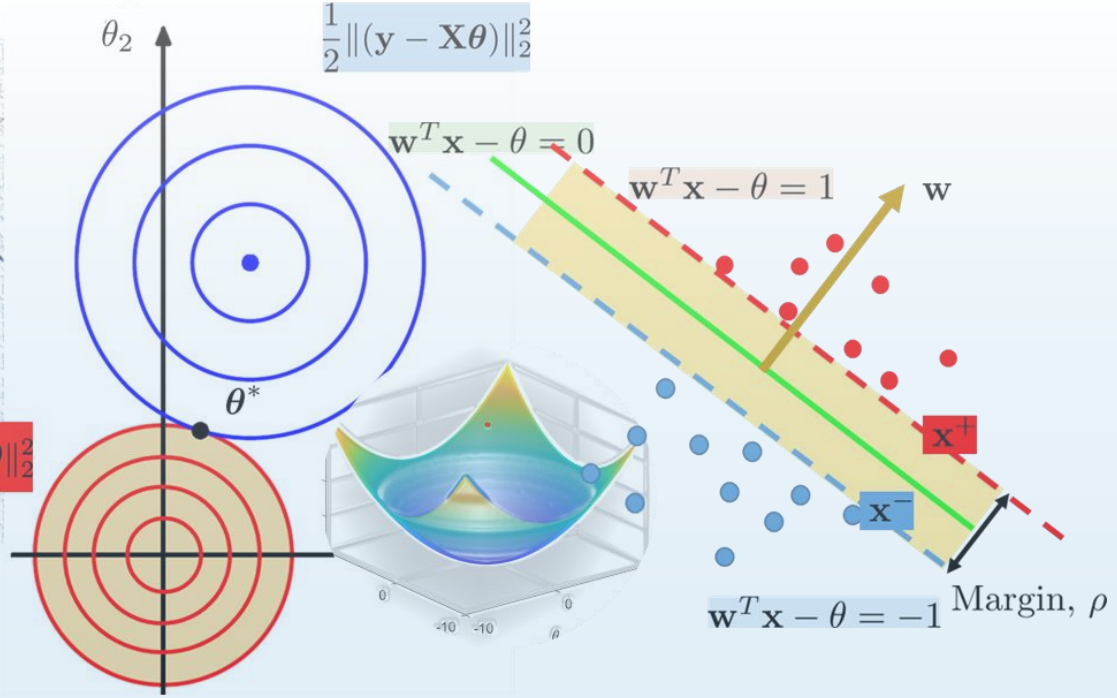


# AI-501 Mathematics for AI

## Vectors – Notation and Basic Operations

Zubair Khalid
School of Science and Engineering

https://www.zubairkhalid.org/ai501_2024.html

LUMS
A Not-for-Profit University

# Outline

- Vectors notation
- Applications
- Basic vector operations

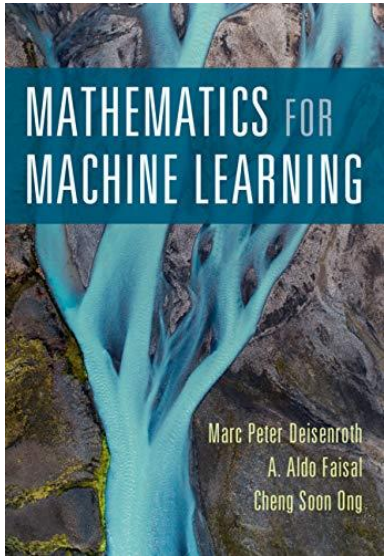


**Chapter 2**

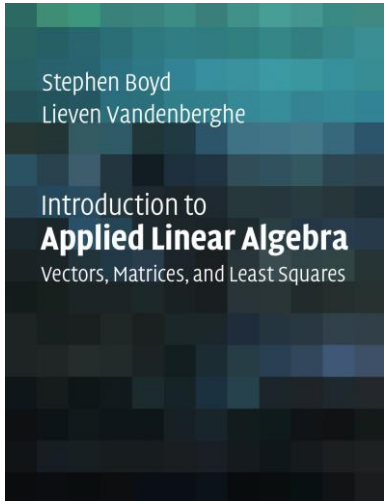


**Chapter 1 and 2**

# Vectors

**Definition:**

A vector is an **ordered** finite list of numbers (real or complex).

**Notation:**

Usually denoted by a letter symbol; stack the list of numbers in an ordered form.

For example, consider a vector of 4 real numbers given by

$$a = \begin{bmatrix} -1.1 \\ 30.1 \\ 6.1 \\ -2.7 \end{bmatrix} \qquad a = \begin{pmatrix} -1.1 \\ 30.1 \\ 6.1 \\ -2.7 \end{pmatrix} \qquad a = (-1.1,\ 30.1,\ 6.1,\ -2.7)$$

**Size of a vector:** Number of elements the vector contains (also referred to as length or dimension).

We usually express vector $b$ of size $n$ as $b \in \mathbf{R}^n$ and call it $n$-vector.

**Entry of a vector:** $b_k$ - $k$-th entry of the vector $b$. For example, $a_2 = 30.1$ for the vector a defined above.

# Vectors

**Zero vector:** A vector with all elements equal to zero.

denoted by $\mathbf{0} \in \mathbf{R}^n$.

**One vector:** A vector with all elements equal to one.

denoted by $\mathbf{1} \in \mathbf{R}^n$.

**Unit vector:** A standard unit vector is vector with all elements zero except one element that is equal to one.

denoted by $e_i \in \mathbf{R}^n$ and is defined as

$$(e_i)_j = \begin{cases} 1 & i = j \\ 0 & i \neq j \end{cases} = \delta_{i,j}$$

**k-Sparse vector:** A vector with at-most **k** non-zero entries.

# Geometric Interpretation

**2-vector:**

# Examples of Vectors - Applications

**Location**

**Displacement, Velocity, Acceleration**

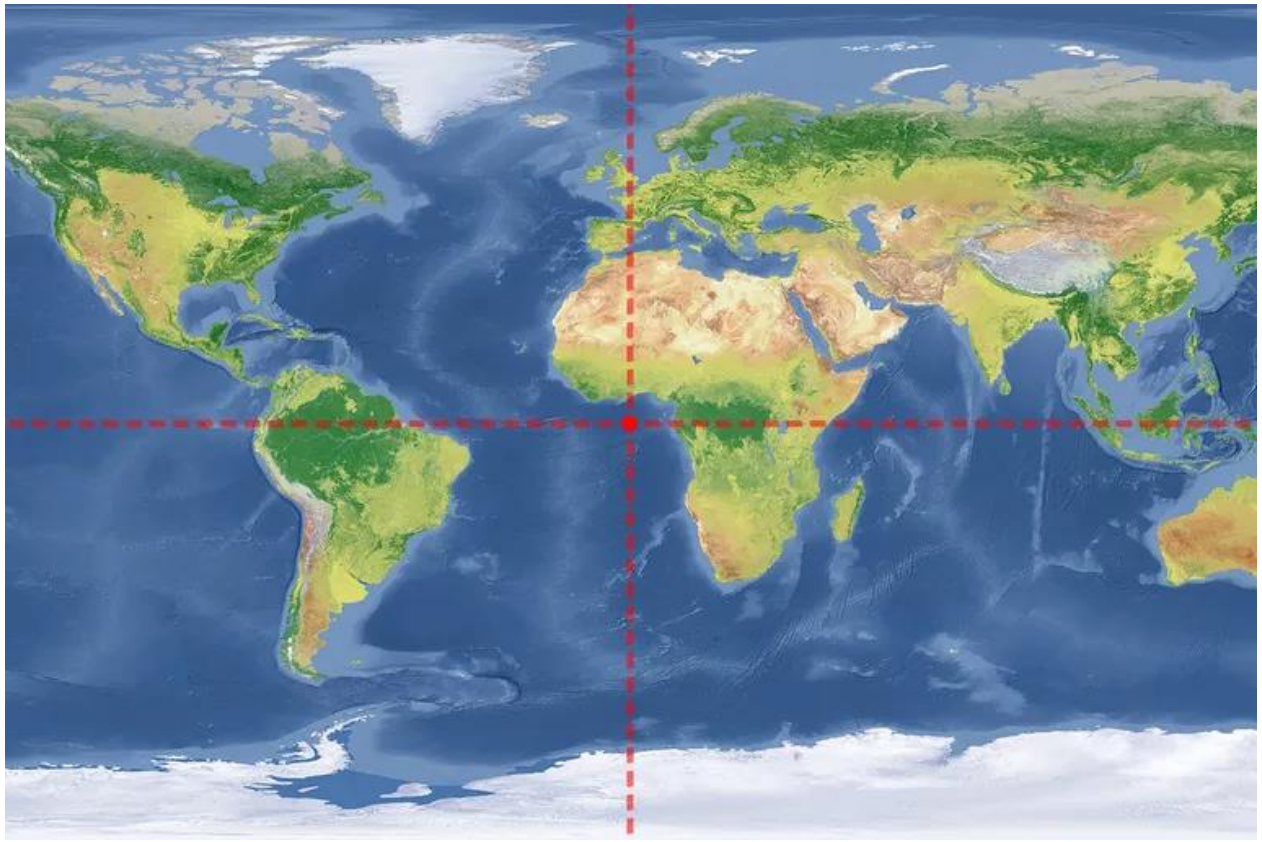

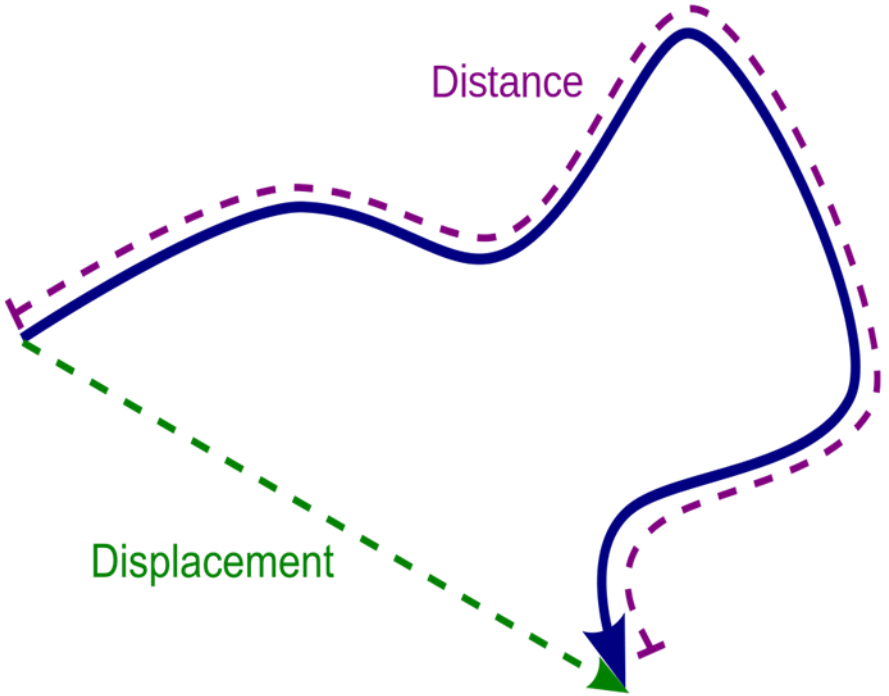

# Examples of Vectors - Applications

**Color**

Each color is represented by 3-vector.

**Quantities**

An n-vector q can represent the amounts or quantities of n different resources or products held (or produced, or required) by an entity such as a company.

For example, n-vector represents the quantity of n products stocked in a warehouse.

**Values across a population**

An n-vector can give the values of some quantity across a population of individuals or entities.

For example, an n-vector **a** can represent the blood pressure of a collection of n patients, with $a_i$ the blood pressure of patient i, for $i = 1, 2, \ldots, n$.

# Examples of Vectors - Applications

**Image**

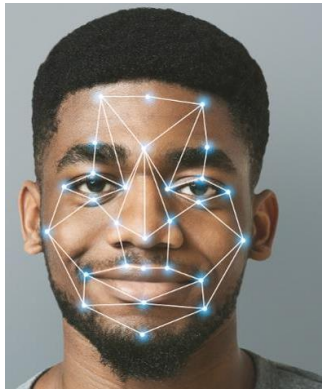

**Time series**

12-vector can represent the average monthly temperature, rainfall, pressure etc of Lahore.

30-vector can represent the number of expected COVID-19 patients in Pakistan cases over the next 30 days.

Other examples include exchange rate, audio, and, in fact, any quantity that varies over time.

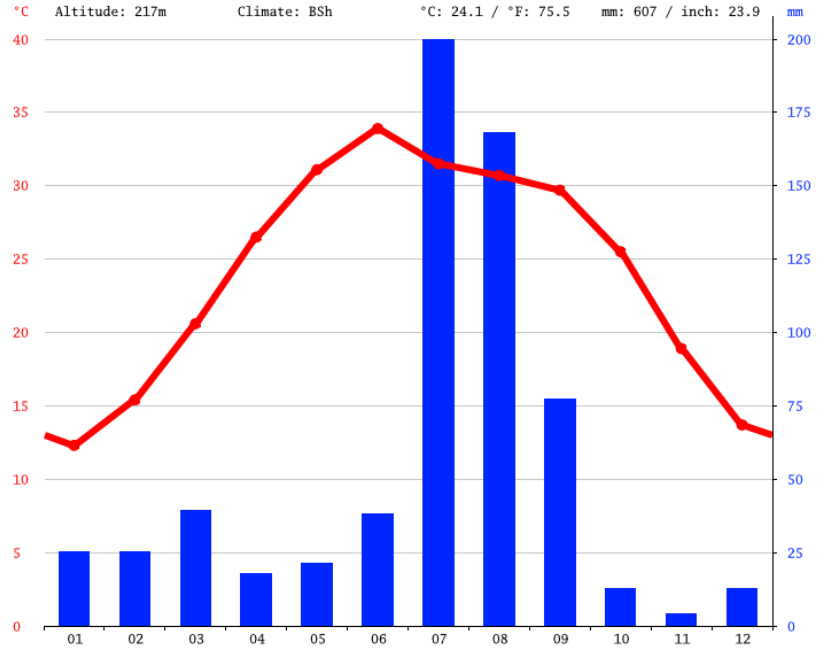

# Examples of Vectors - Applications

**Audio**

Sound (Flute) versus time

8 second sound = 44100x8-vector

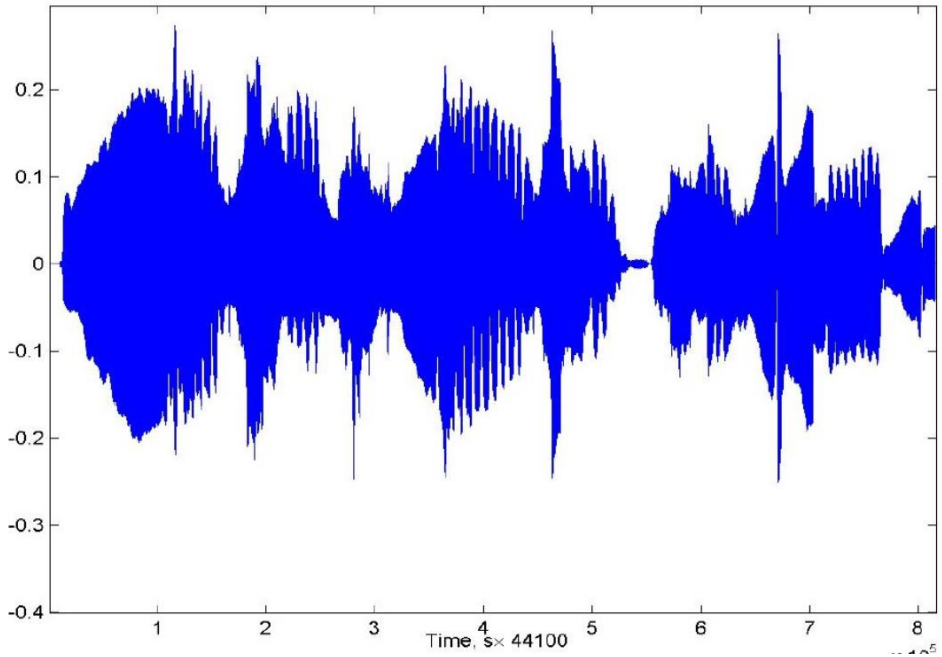


**Feature or Attribute**

*In Machine Learning, classification is mostly carried out collecting features or characteristics derived from the object.*

*Such a vector is sometimes called a feature vector, and its entries are called the features or attributes.*

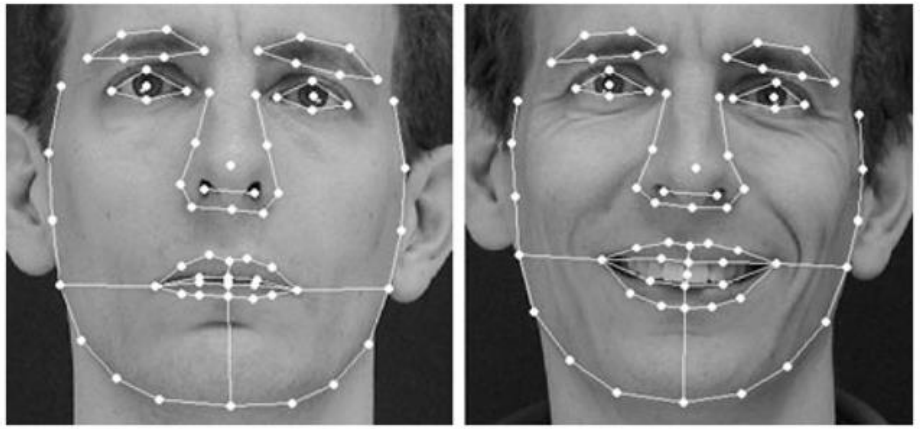

Emotion Recognition