Machine Learning: Introduction

Hao Tang

January 17, 2024

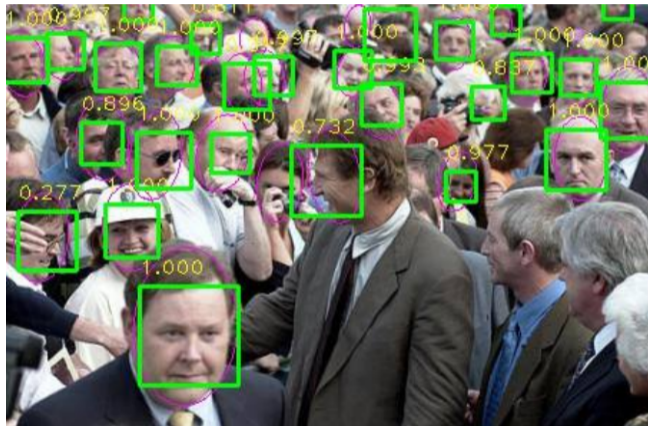`https://homepages.inf.ed.ac.uk/htang2/mlg2023/`

Face recognition

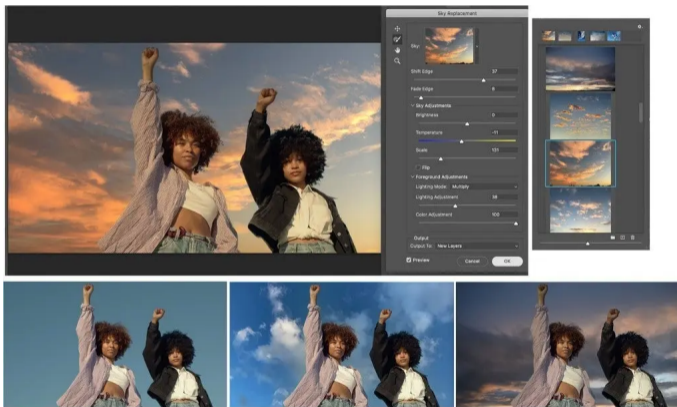Image credit: (Triantafyllidou and Tefas, 2016)

Photo editing

Image credit: https://blog.adobe.com/en/publish/2020/10/20/photoshop-the-worlds-most-advanced-ai-application-for-creatives

Photo editing

Image credit: https://blog.adobe.com/en/publish/2020/10/20/photoshop-the-worlds-most-advanced-ai-application-for-creatives

Video editing

| Source | Driver | Ours | Driver | Ours |

Image credit: https://neeek2303.github.io/MegaPortraits/

AlphaGo

Recommender systems

Speech recognition

Image credit: https://www.bbc.co.uk/news/technology-47893082

Speech verification

Menu | HSBC | Log on

Voice ID
Bye bye passwords

BANK OF SCOTLAND
Branch Finder | Help Centre | Register | Login

Home | Contact us | Call us | Voice ID

Voice ID FAQs

Voice ID - fast, easy and safe

What is Voice ID?

Voice ID is a security feature available to Telephone Banking customers. It confirms your identity using your voice. If you're eligible and hold a bank, savings or credit card account with us, you can set up Voice ID.

BARCLAYS

Log In | Search | Menu

How do I register for Barclays Voice Security and can I do it in branch?

You can only register for Barclays Voice Security over the phone. During your phone conversation with us we'll ask if you'd like to register, and if you do choose to register we'll then capture a unique digital voice pattern for you. This process is completely free, and you can withdraw from the service at any time.

Once we've captured your voice pattern, you'll be registered straight away. You'll be able to access your accounts using Barclays Voice Security the next time you call Telephone Banking, and every time after that.

Barclays Voice Security is currently only available via Telephone Banking, including the initial conversation in which we record your voice pattern.

Robotic vacuum cleaner

Image credit: https://www.irobot.co.uk/en_GB/roomba.html

Autonomous driving

Image credit: https://www.tesla.com/autopilot

Failures

‘It Happened So Fast’: Inside a Fatal Tesla Autopilot Accident

A 2019 crash in Florida highlights how gaps in Tesla's driver-assistance system and distractions can have tragic consequences.

Image credit: `https://www.nytimes.com/2021/08/17/business/tesla-autopilot-accident.html`
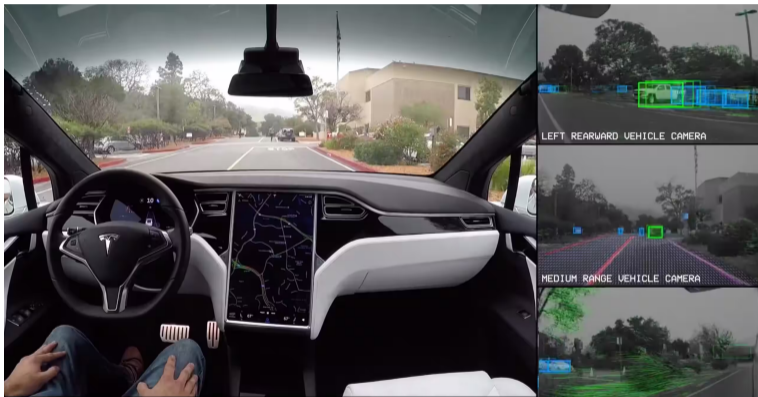
What is machine learning?

What is machine learning?

How would you write a program to recognize hand-written 2s?

What is machine learning?

How would you write a program to recognize hand-written 2s?
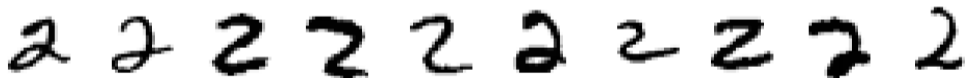
What is machine learning?

How would you write a program to recognize hand-written 2s?

We don't write a program to enumerate all the possible ways of writing 2s.
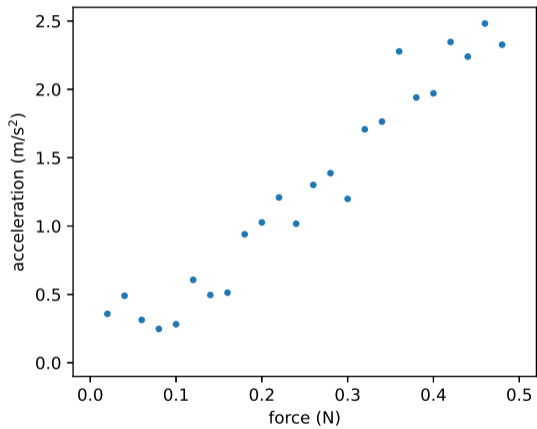We "produce" a program using the examples of 2s.

What is machine learning?
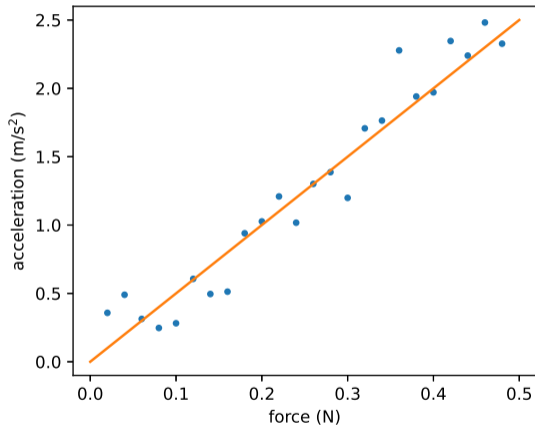
Programming with Data

First example

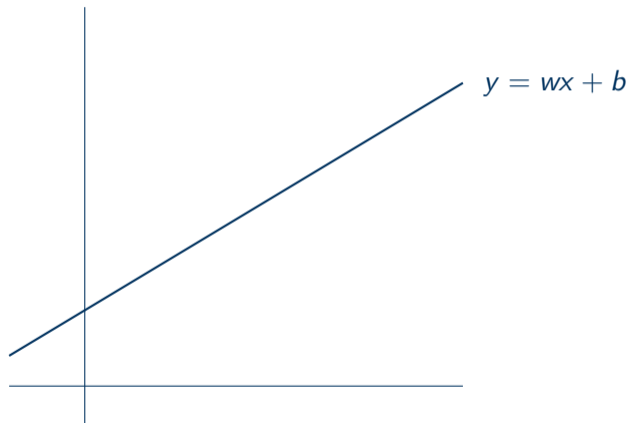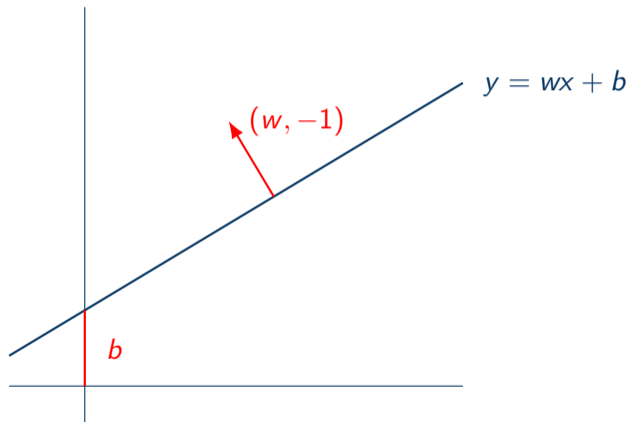| force (N)	acceleration (m/s^2)
0.02	0.358
0.04	0.490
0.06	0.313
0.08	0.247
0.10	0.282
0.12	0.606

First example

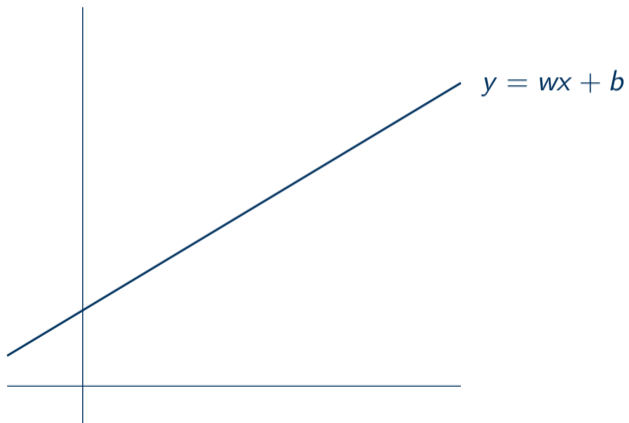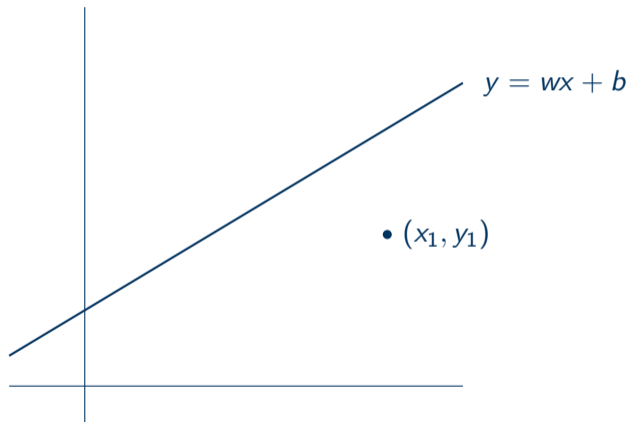# First example

Geometry

$y = wx + b$

Geometry

Geometry

Geometry
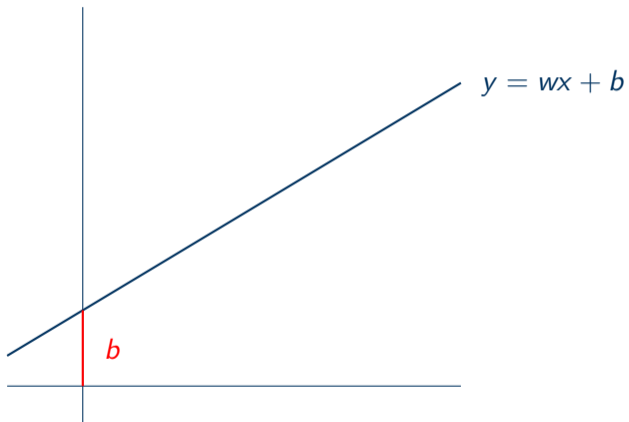
$y = wx + b$

Geometry

$y = wx + b$

$\bullet \, (x_1, y_1)$

Geometry
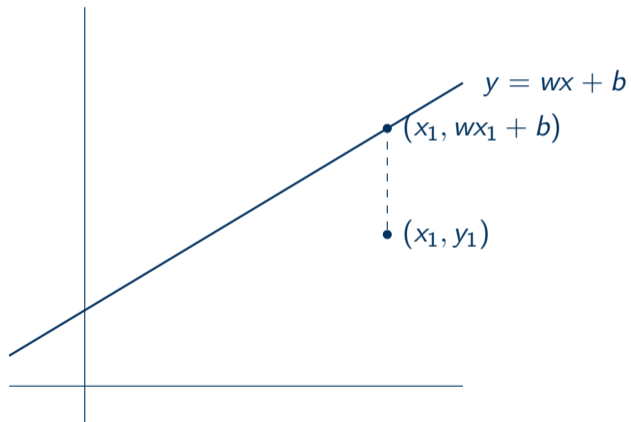
Geometry
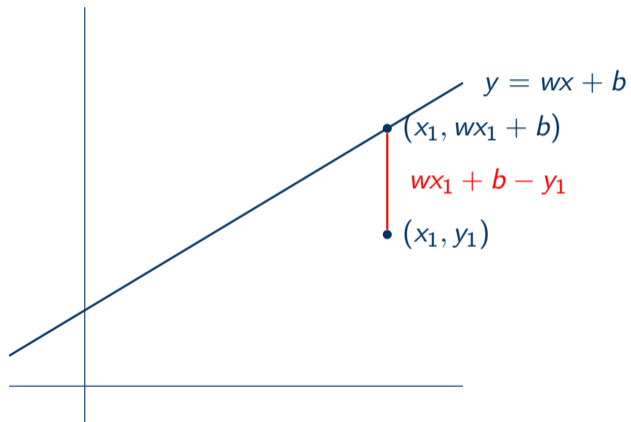
Geometry
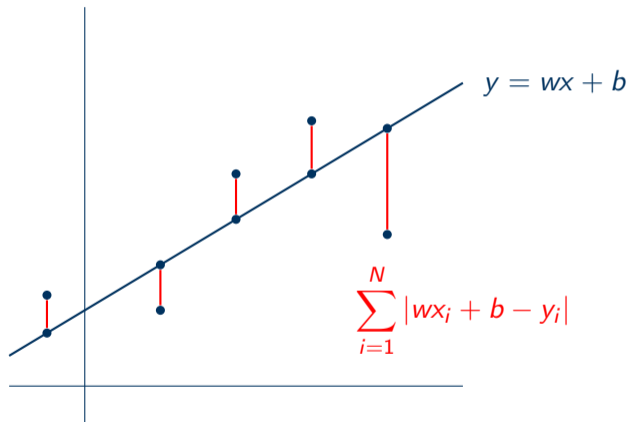
$y = wx + b$

$$\sum_{i=1}^{N} |wx_i + b - y_i|$$

First example
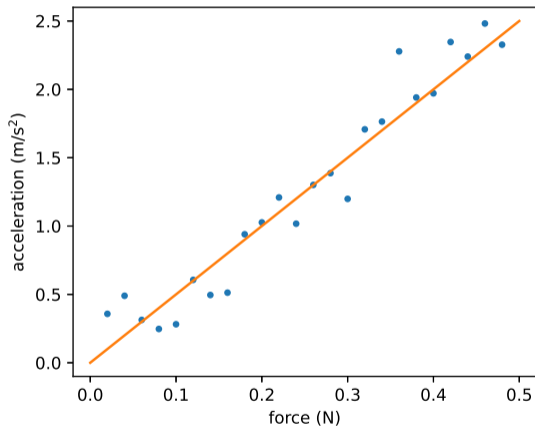
- Given N points $\{(x_1, y_1), (x_2, y_2), \ldots, (x_N, y_N)\}$,

$$\text{sum of absolute errors} \qquad L = \sum_{i=1}^{N} |wx_i + b - y_i|. \qquad (1)$$

- Find w and b that minimizes L.

- Find a function $f(x) = y = wx + b$ that minimizes L.

Terminology

- The N points $\{(x_1, y_1), \ldots, (x_N, y_N)\}$ constitute a **data set**.

Terminology

- The N points $\{(x_1, y_1), \ldots, (x_N, y_N)\}$ constitute a **data set**.

- The value w and b are called **parameters**.

Terminology

- The N points $\{(x_1, y_1), \ldots, (x_N, y_N)\}$ constitute a **data set**.

- The value w and b are called **parameters**.

- The function L is called a **loss function**.

Terminology

- The N points $\{(x_1, y_1), \ldots, (x_N, y_N)\}$ constitute a **data set**.

- The value w and b are called **parameters**.

- The function L is called a **loss function**.

- The act of finding w and b that minimizes L is called **training**.

Terminology

- The N points $\{(x_1, y_1), \ldots, (x_N, y_N)\}$ constitute a **data set**.

- The value w and b are called **parameters**.

- The function L is called a **loss function**.

- The act of finding w and b that minimizes L is called **training**.

- The $x_1, x_2, \ldots, x_N$ are called (input) **features**.

- The $y_1, y_2, \ldots, y_N$ are called (output) **labels**.

Terminology

- The N points $\{(x_1, y_1), \ldots, (x_N, y_N)\}$ constitute a **data set**.

- The value w and b are called **parameters**.

- The function L is called a **loss function**.

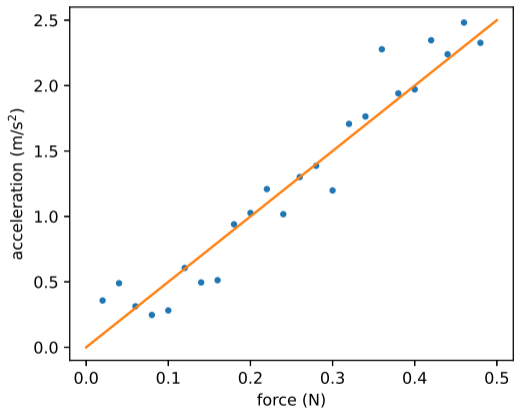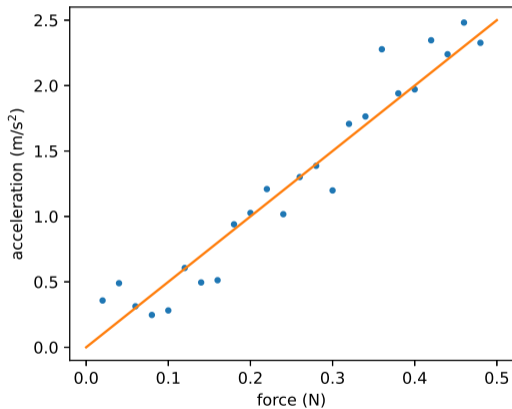- The act of finding w and b that minimizes L is called **training**.

- The $x_1, x_2, \ldots, x_N$ are called (input) **features**.

- The $y_1, y_2, \ldots, y_N$ are called (output) **labels**.

- Specifying the above gives us a **task**.

First example

First example


```python
def acc(f):
    m = 0.2
    return f / m
```

Temperature calibration

ME442 Typical Sensor Connections
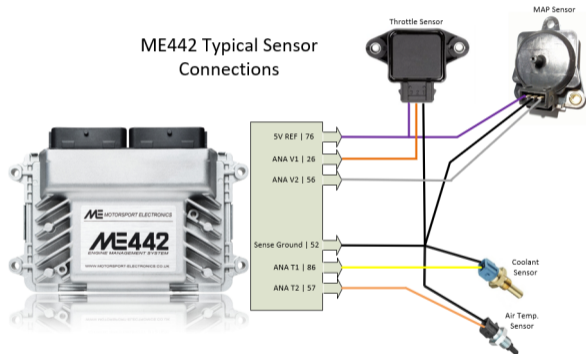
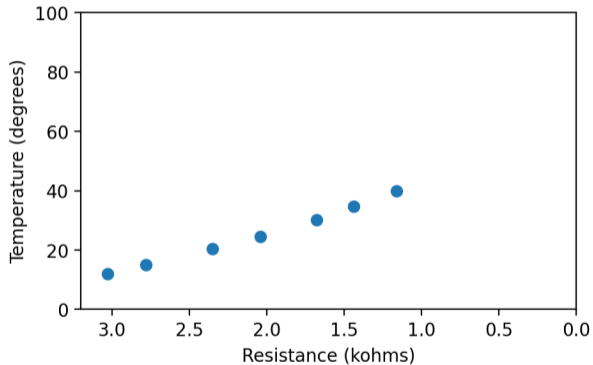Image credit: https://motorsport-electronics.co.uk/onlinehelp/html/Introduction.html

Temperature calibration

Temperature calibration
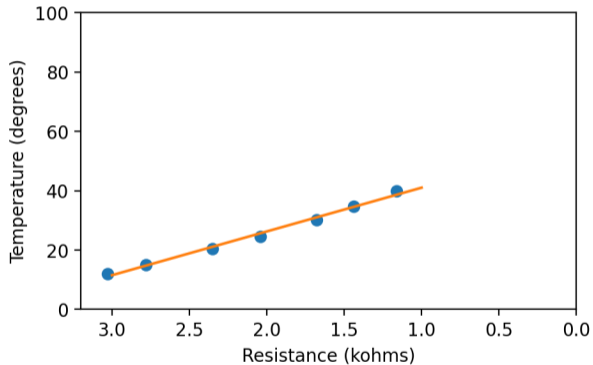
Temperature calibration
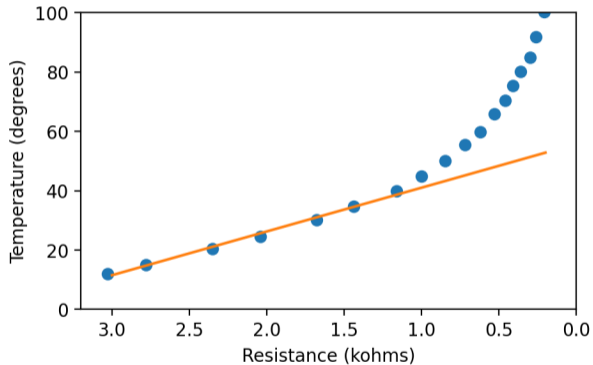
Generalization

- A program is **correct** if it produces the desired output on all input in the input domain.

Generalization

- A program is **correct** if it produces the desired output on all input in the input domain.

- The fact that we use machine learning means that we do not have a good characterization of the input. (If we do, we likely don't need machine learning.)

Generalization

- A program is **correct** if it produces the desired output on all input in the input domain.

- The fact that we use machine learning means that we do not have a good characterization of the input. (If we do, we likely don't need machine learning.)

- Since the input domain is infinitely large, we only develop our program on samples from the input domain.

Generalization

- A program is **correct** if it produces the desired output on all input in the input domain.

- The fact that we use machine learning means that we do not have a good characterization of the input. (If we do, we likely don't need machine learning.)

- Since the input domain is infinitely large, we only develop our program on samples from the input domain.

- A program **generalizes** if it is developed with samples from the input domain but is able to produce the desired output on the entire input domain.

What we will do in this course

Idea Intuition $\longrightarrow$ Math $\longrightarrow$ Implementation Program $\longrightarrow$ Correctness

Things we need

- Calculus

- Linear algebra

- Probability

- python, numpy, matplotlib

- ~~Jupiter notebook~~

Connections to other courses

- Foundations of Data Science (FDS)

- Applied Machine Learning (AML)

- Machine Learning and Pattern Recognition (MLPR)

- Probabilistic Modeling and Reasoning (PMR)

- Machine Learning Practical (MLP)

- Machine Learning Theory (MLT)

- Reinforcement learning (RL)

Connections to other courses

- Foundations of natural language processing (FNLP)

- Accelerated natural language processing (ANLP)

- Natural language understanding, generation, and machine translation (NLU+)

- Speech processing (in PPLS)

- Automatic speech recognition (ASR)

- Speech synthesis (in PPLS)

- Image and vision computing

- Advanced robotics

Logistics

- Course website: `https://homepages.inf.ed.ac.uk/htang2/mlg2023/`

- Textbooks

- Exercises and notes

- Applications

- Other sources
  - `https://twitter.com/PinakiLaskar/status/1329748899347767296`
  - `https://twitter.com/PhDemetri/status/1723802334545813507`

- Content flexibility

- Coming to lectures and tutorials