

Machine Learning

Lecture 1: Introduction

Hao Tang

September 22, 2022

<https://homepages.inf.ed.ac.uk/htang2/mlg2022/>

Face recognition



Image credit: (Triantafyllidou and Tefas, 2016)

Photo editing

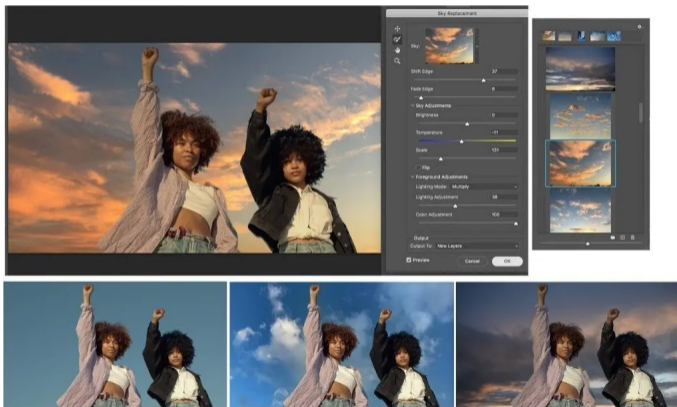


Image credit: <https://blog.adobe.com/en/publish/2020/10/20/photoshop-the-worlds-most-advanced-ai-application-for-creatives>

Photo editing



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

Video editing



Image credit: <https://samsunglabs.github.io/MegaPortraits/>

AlphaGo



Image credit: <https://www.theguardian.com/technology/2016/mar/15/alphago-what-does-google-advanced-software-go-next> (left)
<https://twitter.com/demishassabis/status/692430224382902272> (right)

Recommender systems



NETFLIX



Spotify[®]



TikTok



YouTube

Speech recognition



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

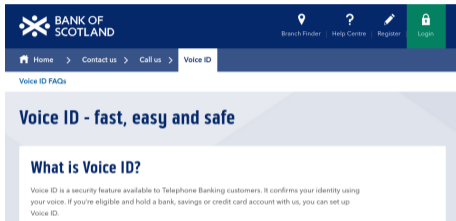
Speech verification




Menu  Log in

Voice ID
Bye bye passwords

A woman with curly hair is smiling while talking on a mobile phone. The background is a blurred indoor setting, possibly a cafe or office.



 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.



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/en_GB/videos/autopilot-self-driving-hardware-neighborhood-short

Failures



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




Geometry




Geometry




Geometry




Geometry




Geometry



Geometry



Geometry



Given N points $\{(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)\}$,

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


The line that minimizes the error would be


$$\hat{w}, \hat{b} = \underset{w, b}{\operatorname{argmin}} L. \quad (2)$$


Optimization

How do we solve

$$\min_{w,b} \sum_{i=1}^N |wx_i + b - y_i|? \quad (3)$$







- The problem is that $|a|$ is not differentiable for all $a \in \mathbb{R}$.
- Instead, we modify the loss function to

$$L = \sum_{i=1}^N (wx_i + b - y_i)^2. \quad (4)$$

- The plan is to solve

$$\frac{\partial}{\partial w} L = 0 \quad \frac{\partial}{\partial b} L = 0 \quad (5)$$

$$\frac{\partial}{\partial b} \sum_{i=1}^N (wx_i + b - y_i)^2 = \sum_{i=1}^N 2(wx_i + b - y_i) \quad (6)$$

$$= 2 \left(w \sum_{i=1}^N x_i + nb - \sum_{i=1}^N y_i \right) = 0 \quad (7)$$

$$b = \frac{1}{n} \sum_{i=1}^N y_i - w \frac{1}{n} \sum_{i=1}^N x_i = \bar{y} - w\bar{x} \quad (8)$$


$$L = \sum_{i=1}^N (w(x_i - \bar{x}) - (y_i - \bar{y}))^2 \quad (9)$$

$$\frac{\partial}{\partial w} L = \sum_{i=1}^N 2(w(x_i - \bar{x}) - (y_i - \bar{y}))(x_i - \bar{x}) \quad (10)$$

$$= 2 \left(w \sum_{i=1}^N (x_i - \bar{x})(x_i - \bar{x}) - \sum_{i=1}^N (y_i - \bar{y})(x_i - \bar{x}) \right) = 0 \quad (11)$$

$$w = \frac{\sum_{i=1}^N (y_i - \bar{y})(x_i - \bar{x})}{\sum_{i=1}^N (x_i - \bar{x})(x_i - \bar{x})} \quad (12)$$

First example



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

Temperature calibration



ME442 Typical Sensor Connections



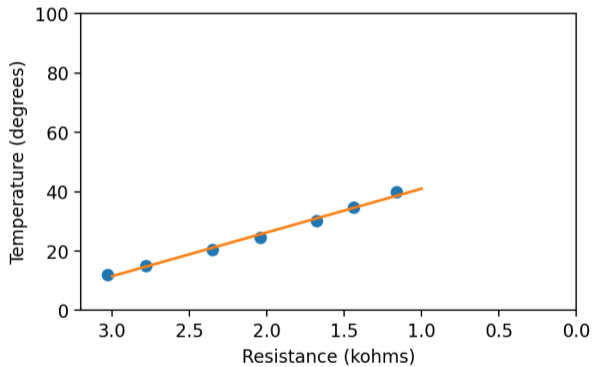


Image credit: <https://motorsport-electronics.co.uk/onlinehelp/html/Introduction.html>


Temperature calibration



Temperature calibration



Temperature calibration



Generalization

- Machine learning models are tools for **prediction**.
- The ability to predict is called **generalization**.
- A learned program is “correct” if it is able to generalize.

Recipe

- Choose a **task**, e.g., classification or regression. (Lecture 5)
- Choose a **data set**. (Lecture 28, 29)
- Choose a **loss**, e.g., mean squared error or cross entropy. (Lecture 4, 5, 6)
- Choose a **model**, e.g., two-layer neural network. (Most of this course)
- Choose an **optimization algorithm**, e.g., stochastic gradient descent. (Lecture 8, 9, 11)
- Check **generalization**. (Lecture 15, 16, 17, 18)

Things we need

- Calculus
- Linear algebra
- Probability
- python, numpy, matplotlib
- ~~Jupyter notebook~~

What you will not learn in this course

What you will not learn in this course

- ~~Use sklearn~~

What you will not learn in this course

- ~~Use sklearn~~
- ~~Stanford CS229: Machine Learning by Andrew Ng~~

What you will not learn in this course

- ~~Use sklearn~~
- ~~Stanford CS229: Machine Learning by Andrew Ng~~
- ~~Kaggle~~

What you will learn in this course

- Formalize a ML problem with math
- Read and understand ML theorems
- Turn ML algorithms into programs

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/mlg2022/>
- No textbook. Optional reading on the website.
- 30 hours of lectures
 - Mon at 15:10
 - Thu at 15:10
 - Fri at 15:10
- TA drop-in session Wednesday at 12pm


Logistics

- 5 tutorials on Friday
- 1 practice exam
- 1 coursework (20%)
- 1 final exam (80%)

Lecturers



Kia Nazarpour



Hao Tang



Hiroshi Shimodaira

TAs

Borislav Ikonov



Sung-Lin Yeh



Ramon Sanabria