

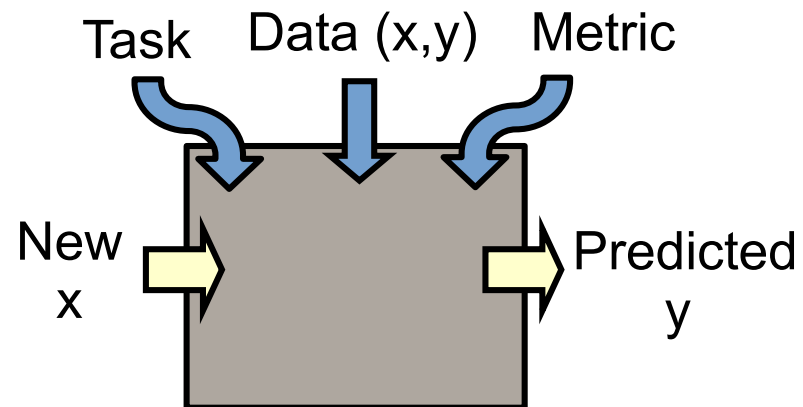
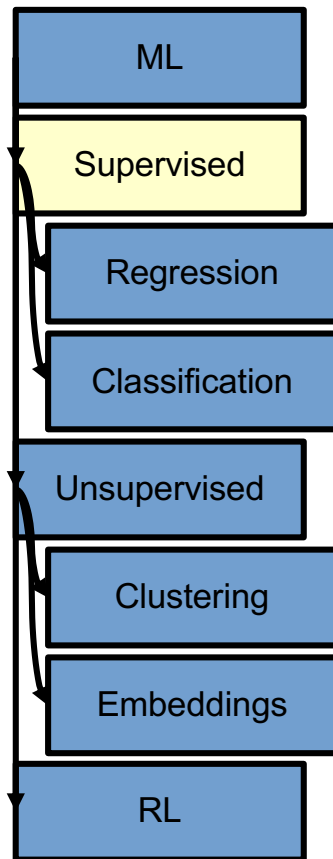
CS181: Introduction to Machine Learning

Lecture 4 (Linear classification)



Spring 2021


Finale Doshi-Velez and David C. Parkes
Harvard Computer Science

Machine Learning Taxonomy



Terminology: Classification





+ graphical models, reinforcement learning



Example: Digit recognition for mail sorting

3 6 8 1 7 9 6 6 9 1
6 7 5 7 8 6 3 4 8 5
2 1 7 9 7 1 2 8 4 5
4 8 1 9 0 1 8 8 9 4
7 6 1 8 6 4 1 5 6 0
7 5 9 2 6 5 8 1 9 7
2 2 2 2 2 3 4 4 8 0
0 2 3 8 0 7 3 8 5 7
0 1 4 6 4 6 0 2 4 3
7 1 2 8 7 6 9 8 6 1


The USPS is considered something of a pioneer in the field of machine learning, and they were one of the first organizations to start making substantial investments in the technology. After researching options for over a decade, they deployed their first computer prototype capable of reading handwritten text in 1997 with the help of the [University of Buffalo's Center for Excellence in Document Analysis and Recognition \(CEDAR\)](#). The prototype correctly identified the addresses on only 10% of envelopes it read with a 2% error rate. Today, machines at the USPS process approximately [98% of hand-addressed mail and 99.5% of type-written addresses](#).

https://www.enterpriseai.news/solution_content/hpe/governmentacademia/machine-learning-applications-for-the-modern-enterprise/

Example: Citizen Science



Welinder et al., 2010



DATA MATCH

00

MONTHS

02

DAYS

15

HOURS

53

MINS

38

SECS

Example: Facebook groups

The Guardian, 2021

Facebook

Kari Paul

Thu 4 Feb 2021 11:00 GMT



32

'It let white supremacists organize': the toxic legacy of Facebook's Groups




▲ By the time Facebook banned QAnon content in 2020, a Guardian report had exposed that groups dedicated to the dangerous conspiracy theory were spreading on the platform at a rapid pace. Photograph: Carlos Barria/Reuters

Facebook has said it will no longer algorithmically recommend political groups to users, but experts warn that isn't enough

Advertisement






PRML, Bishop

Red = +1


Blue = -1




PRML, Bishop

Red = +1

Blue = -1



PRML, Bishop
Red = +1
Blue = -1



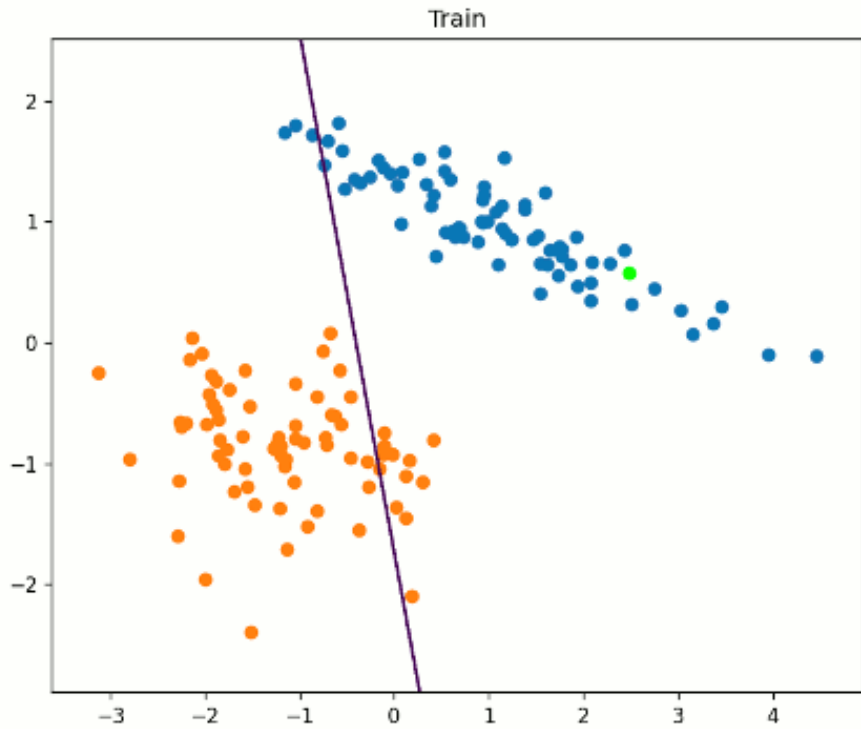
PRML, Bishop

Red = +1

Blue = -1


Separable

Iteration: 1/2; Point: 1/150

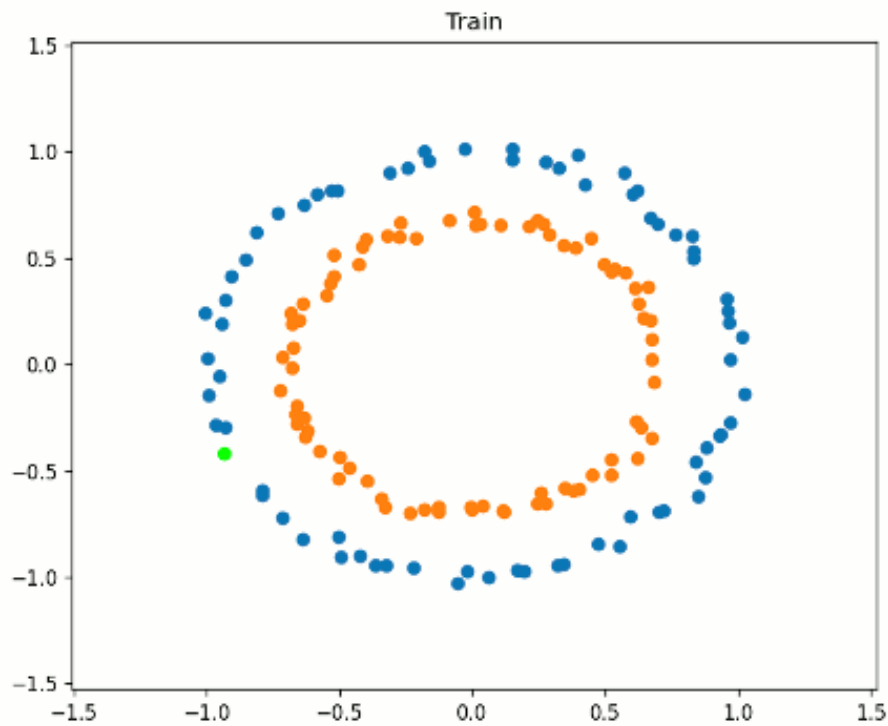


Non-separable

Iteration: 1/100



Iteration: 1/3; Point: 1/150




<https://towardsdatascience.com/perceptron-explanation-implementation-and-a-visual-example-3c8e76b4e2d1>

Use a nonlinear basis

$$\phi(x) = \begin{bmatrix} x_1 \\ x_2 \\ x_1^2 \\ x_1x_2 \\ x_2^2 \end{bmatrix}$$

Show projection of classifier back to 2D space




Perceptron
(and hinge loss,
more
generally) can't
choose
between
multiple
separators




Frank Rosenblatt 1928–1969

Rosenblatt's perceptron played an important role in the history of machine learning. Initially, Rosenblatt simulated the perceptron on an IBM 704 computer at Cornell in 1957, but by the early 1960s he had built


Mark 1 Perceptron Hardware



Inputs via camera system (20x20)



Patch board, to reconfigure inputs



One of the racks of adaptive weights, each implemented using a rotary variable resistor, driven by an electric motor