



ML Framework: From Lines to Perceptrons

Dr. Fayyaz Minhas

Department of Computer Science
University of Warwick

<https://warwick.ac.uk/fac/sci/dcs/teaching/material/cs909/>


Intuition of Linear Discriminants

- Paper Cutting

Fold and Cut Theorem Video:

<https://www.youtube.com/watch?v=ZREp1mAPKTM>

- How can we solve non-linear classification?
 - By folding the space on which examples lie and then making a single straight cut
 - Notice how folding changes the distance between points
 - How to achieve such folding?
 - One way is to transform the data



<https://www.youtube.com/watch?v=ZREp1mAPKTM>

The Fold-and-Cut Theorem implies that any pattern can be achieved with a single straight cut if the paper (or space) is folded appropriately.

Thus, it is theoretically possible to partition any space into regions containing positive and negative training examples no matter how complex such a boundary is by simply folding the feature space appropriately and using a linear classifier (single straight cut).

https://en.wikipedia.org/wiki/Fold-and-cut_theorem

Preliminaries and Intuition

Preliminaries

- Equations of lines and their properties

$$w_1x_1 + w_2x_2 + b = 0$$

$$x_2 = \frac{-w_1}{w_2}x_1 + \frac{-b}{w_2}$$

$$x_2 = mx_1 + c$$

$$f(\mathbf{x}; \mathbf{w}) = w_1x_1 + w_2x_2 + b = 0$$

$$f(\mathbf{x}; \mathbf{w}) = [w_0 \quad w_1 \quad w_2] \begin{bmatrix} 1 \\ x_1 \\ x_2 \end{bmatrix} = \mathbf{w}^T \mathbf{x} = 0$$


If a point $x = (x_1, x_2)$ is on the line then $f(\mathbf{x}; \mathbf{w}) = w_1x_1 + w_2x_2 + b = 0$

If it is above the line then $f(\mathbf{x}; \mathbf{w}) = w_1x_1 + w_2x_2 + b > 0$

If it is below the line then $f(\mathbf{x}; \mathbf{w}) = w_1x_1 + w_2x_2 + b < 0$



Line Equation: $w_1x_1 + w_2x_2 + b = 0$



<https://foxtrotmike.github.io/CS909/lines.html>

Preliminaries

- Distance function

$$d(\mathbf{a}, \mathbf{b}) = \|\mathbf{a} - \mathbf{b}\| = \sqrt{(\mathbf{a}_1 - \mathbf{b}_1)^2 + (\mathbf{a}_2 - \mathbf{b}_2)^2}$$

- Relation with Dot Product

$$d^2(\mathbf{a}, \mathbf{b}) = \|\mathbf{a} - \mathbf{b}\|^2 = (\mathbf{a}_1 - \mathbf{b}_1)^2 + (\mathbf{a}_2 - \mathbf{b}_2)^2 = \mathbf{a}^T \mathbf{a} + \mathbf{b}^T \mathbf{b} - 2\mathbf{a}^T \mathbf{b}$$

- If \mathbf{a}, \mathbf{b} are unit vectors (i.e., $\mathbf{a}^T \mathbf{a} = \|\mathbf{a}\|^2 = \mathbf{b}^T \mathbf{b} = \|\mathbf{b}\|^2 = 1$) then: $d^2(\mathbf{a}, \mathbf{b}) = 2 - 2(\mathbf{a}^T \mathbf{b})$
- Or the farther away or different two points are, the lower their dot product and vice-versa**
- We can also have more generalized dot products called kernel functions that can measure similarity between two objects in a different way**
 - Linear kernel: $k(\mathbf{a}, \mathbf{b}) = \mathbf{a}^T \mathbf{b}$
 - Polynomial kernel: $k(\mathbf{a}, \mathbf{b}) = (\mathbf{a}^T \mathbf{b})^2$
 - Gaussian or Radial Basis Function (RBF) Kernel: $k(\mathbf{a}, \mathbf{b}) = \exp(-\lambda \|\mathbf{a} - \mathbf{b}\|^2)$
 - Mahalanobis Kernel: $k(\mathbf{a}, \mathbf{b}; \mathbf{M}) = \exp(-(\mathbf{a} - \mathbf{b})^T \mathbf{M} (\mathbf{a} - \mathbf{b}))$
 - Exponential kernel: $k(\mathbf{a}, \mathbf{b}; \mathbf{W}_a, \mathbf{W}_b) = \exp\left(\frac{1}{\sqrt{d}} \langle \mathbf{a} \mathbf{W}_a, \mathbf{b} \mathbf{W}_b \rangle\right)$
 - Asymmetric, non-mercer kernels

Angle of Vector A (degrees):


Angle of Vector B (degrees):

Distance: 0.00

Dot Product: 1.00

Polynomial Kernel: 1.00

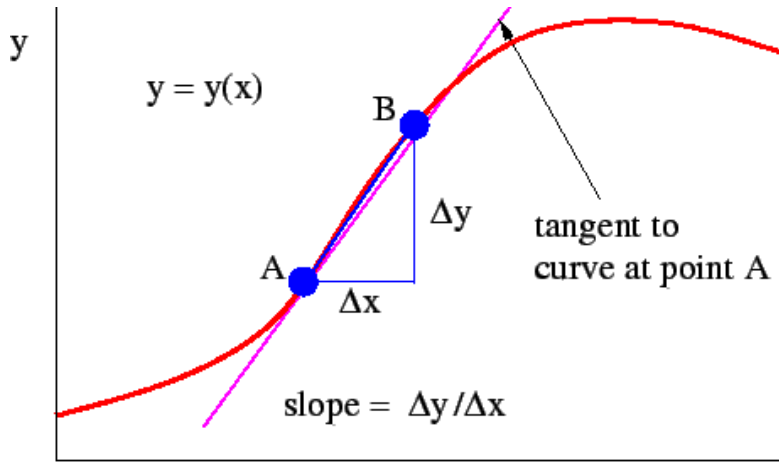

RBF Kernel: 1.00




https://foxtrotmike.github.io/CS909/distance_dot.html

Preliminaries

- Gradients



$$\nabla f(\mathbf{x}) = \begin{bmatrix} \frac{\partial f(\mathbf{x})}{\partial x^{(1)}} \\ \frac{\partial f(\mathbf{x})}{\partial x^{(2)}} \\ \vdots \\ \frac{\partial f(\mathbf{x})}{\partial x^{(d)}} \end{bmatrix}$$



$$f(x, y) = x^2 + y^2 \text{ and } \nabla f(x, y) = \begin{bmatrix} \frac{\partial f(x, y)}{\partial x} \\ \frac{\partial f(x, y)}{\partial y} \end{bmatrix} = \begin{bmatrix} 2x \\ 2y \end{bmatrix}$$

<https://en.wikipedia.org/wiki/Gradient>

<https://foxtrotmike.github.io/CS909/gradients.html>

Preliminaries: Finding minima and maxima of functions

- Given a function $f(w)$
- Take the derivative
- Substitute the derivative to zero
- Solve for x when $\frac{df}{dw} = 0$
- Works when we can solve for w



$$\begin{aligned}f(w) &= (w - 0.5)^2 \\ \frac{df}{dw} &= 2(w - 0.5) = 0 \\ w^* &= 0.5\end{aligned}$$

$$\begin{aligned}f(w) &= (w - 0.5)^2 + \sin(4w) \\ \frac{df}{dw} &= 2(w - 0.5) + 4\cos(4w) = 0 \\ w^* &=?\end{aligned}$$

Preliminaries: Gradient Descent

- In order to find the minima of a function, keep taking steps along a direction opposite to the gradient of the function

$$\mathbf{w}^{(k+1)} = \mathbf{w}^{(k)} - \alpha \nabla f(\mathbf{w}^{(k)})$$



GD Implementation

```
import numpy as np

def gd(fxn,dfxn,w0=0.0,lr = 0.01,eps=1e-4,nmax=1000, history = True):
    """
    Implementation of a gradient descent solver.
    fxn: function returns value of the target function for a given w
    dfxn: gradient function returns the gradient of fxn at w
    w0: initial position [Default 0.0]
    lr: learning rate [0.001]
    eps: min step size threshold [1e-4]
    nmax: maximum number of iters [1000]
    history: whether to store history of x or not [True]
    Returns:
    w: argmin_x f(w)
    converged: True if the final step size is less than eps else false
    H: history
    """
    H = []
    w = w0
    if history:
        H = [[w,fxn(w)]]
    for i in range(nmax):
        dw = -lr*dfxn(w) #gradient step
        if np.linalg.norm(dw)<eps: # we have converged
            break
        if history:
            H.append([w+dw,fxn(w+dw)])
        w = w+dw #gradient update
    converged = np.linalg.norm(dw)<eps
    return w,converged,np.array(H)
```

```
if __name__=='__main__':
    import matplotlib.pyplot as plt
    def myfunction(w):
        z = (w-0.5)**2#+np.sin(4*w)
        return z
    def mygradient(w):
        dz = 2*(w-0.5)#+4*np.cos(4*w)
        return dz

    wrange = np.linspace(-3,3,100)
    #select random initial point in the range
    w0 = np.min(wrange)+(np.max(wrange)-np.min(wrange))*np.random.rand()

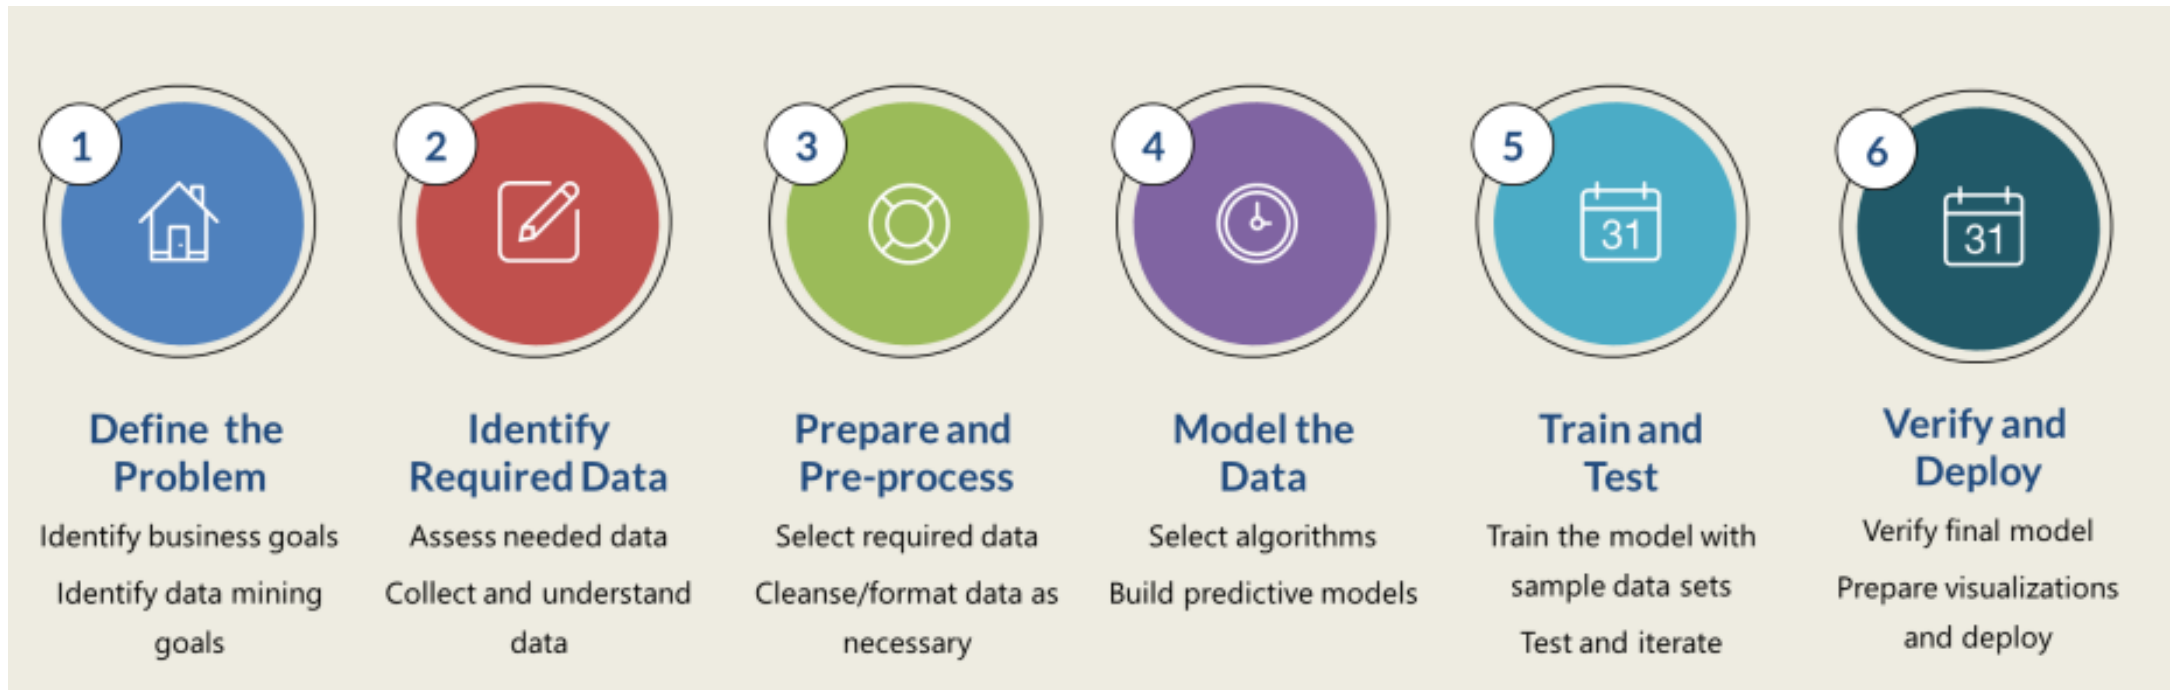
    w,c,H = gd(myfunction,mygradient,w0=w0,lr = 0.01,eps=1e-4,nmax=1000, history = True)

    plt.plot(wrange,myfunction(wrange)); plt.plot(wrange,mygradient(wrange));
    plt.legend(['f(w)', 'df(w)'])
    plt.xlabel('w');plt.ylabel('value')
    s = 'Convergence in '+str(len(H))+' steps'
    if not c:
        s = 'No '+s
    plt.title(s)
    plt.plot(H[0,0],H[0,1], 'ko', markersize=10)
    plt.plot(H[:,0],H[:,1], 'r.-')
    plt.plot(H[-1,0],H[-1,1], 'k*', markersize=10)
    plt.grid(); plt.show()
```

<https://github.com/foxtrotmike/CS909/blob/master/gd.py>


https://github.com/foxtrotmike/CS909/blob/master/dm_lab_2_fm.ipynb

Steps in the development of a data science model



Constructs of a Data Mining System for Prediction

- Identify the objective
 - Identify the unit of classification (example)
 - Image block, protein sequence,





Classification

- Given a set of data points (also called examples)
 - In solving a classification problem, the first step is to identify the unit of classification or “what is an example”
- Such that each example is represented by a feature vector
 - Representation of the example in terms of feature vector
- Assign a class label to each example
 - Such as apple, orange or mango

- Training Data
 - Set of examples for which both feature vectors and labels are available for “tuning”
 - Finding a mathematical function (called a classifier) that can be used to assign these labels

- Validation Data
 - Set of examples (with known labels) that are used to evaluate how well the trained classifier is expected to generalize to novel cases

- Test Data
 - Data for which the labels are not known and the ML model is used to find their labels



Classification

- The Objective of Classification is to assign class labels $y \in \{c_1, c_2, \dots, c_M\}$



- to a given feature vector $x = \begin{bmatrix} x^{(1)} \\ x^{(2)} \\ \vdots \\ x^{(d)} \end{bmatrix}$

- The classifier may use previously known and available training data

– Good generalization, Good memorization


- The training data comprises of classified data points:

$$X = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}, \mathbf{x}_i^{(d \times 1)} = \begin{bmatrix} x_i^{(1)} \\ x_i^{(2)} \\ \vdots \\ x_i^{(d)} \end{bmatrix}$$
$$Y = \{y_1, y_2, \dots, y_N\}, y_i \in \{c_1, c_2, \dots, c_M\}$$



Classification Approaches: Nearest Neighbor and kNN

$$D(\mathbf{x}_a, \mathbf{x}_b) = \sqrt{\left(x_a^{(1)} - x_b^{(1)}\right)^2 + \left(x_a^{(2)} - x_b^{(2)}\right)^2}$$



- Python Warm-up Lab Exercise
- https://github.com/foxtrotmike/CS909/blob/master/DM_1_kNN.ipynb

Example (k=1)-Nearest Neighbor Classification

K-Nearest Neighbors Demo

Instructions:

Select the type of point you want to place (Red or Blue) using the dropdown menu.

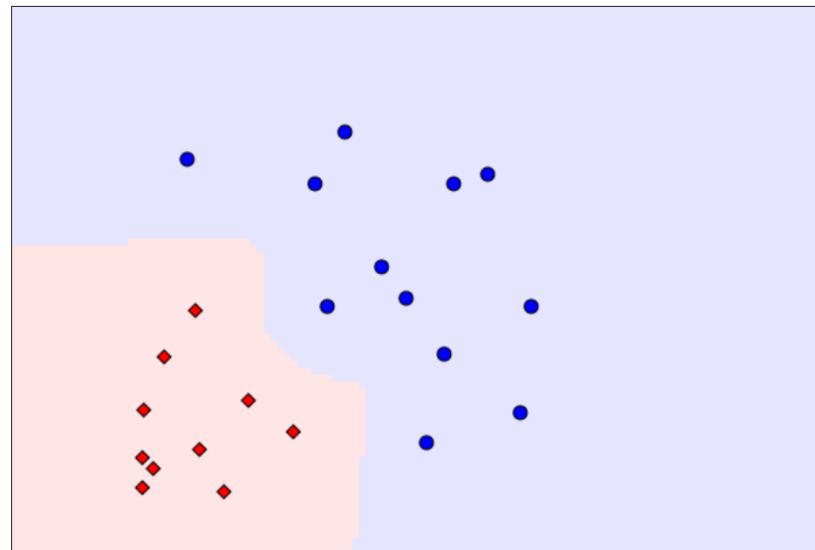
Click anywhere on the canvas to place the selected point.

Click on an existing point to select it. The 'K' nearest neighbors of the selected point will be highlighted with lines.

Adjust the number of neighbors (K) using the input box to see how the neighbors and decision boundary change dynamically.

Select a distance metric from the dropdown to observe its effect on the decision boundaries.

Use the "Clear Points" button to reset the canvas and start over.



Point Type: Num Neighbors (K): Distance Metric:

(c) Fayyaz Minhas


What are some issues with the k-NN classifier?

- Class imbalance.
- In higher-dimensions, there is essentially no notion of distance.

Demo: <https://foxtrotmike.github.io/CS909/knn.html>


Classification Approaches: Supervised...

- Nonlinear Classification boundary




Linear Separability

- If data points can be separated by a linear discriminant then that dataset/classification problem is called “linearly separable”
- Mathematically,
 - If there exists a linear function
$$f(\mathbf{x}; \mathbf{w}) = w_1x^{(1)} + w_2x^{(2)} + \dots + w_dx^{(d)} + b = 0$$
such that
 - If $y_i = +1$, then $f(\mathbf{x}_i; \mathbf{w}) > 0$
 - And If $y_i = -1$, then $f(\mathbf{x}_i; \mathbf{w}) < 0$




Classification Approaches: Supervised...

- **Generalization and Memorization**
 - Remembering everything is not learning
 - The true test of learning is handling similar but unseen cases, i.e., Generalization



Has great memorization but may generalize poorly (under certain assumptions)



Has lesser memorization but may generalize better (under certain assumptions)

Discriminant based classification

- In this type of classification, the objective is to learn a function or, in the case of more than 2 classes, a set of functions from training data which can generate decisions for test data such that the classes in the data can be separated

- Class label assignment: $c(\mathbf{x}) = \operatorname{argmax}_{k=1,\dots,M} f_k(\mathbf{x})$

- $f_k(\mathbf{x})$ tells you the ‘ k -classiness’ of an example \mathbf{x}

- If $M = 2$


- Choose class-1 if $f_1(\mathbf{x}) \geq f_2(\mathbf{x})$, i.e., $f_1(\mathbf{x}) - f_2(\mathbf{x}) \geq 0$
- Otherwise assign it to class-2, i.e., $f_1(\mathbf{x}) - f_2(\mathbf{x}) < 0$
- We can thus replace the two functions with a single function

$$f(\mathbf{x}) = f_1(\mathbf{x}) - f_2(\mathbf{x})$$




Assign to positive class if $f(\mathbf{x}) \geq 0$, otherwise negative

$f(\mathbf{x}) = 0$ separates the two classes and is called the discriminant





If the function(s) are linear, the classifier is called a linear discriminant




Application: Computational Pathology



Block preparation
Addition of dyes

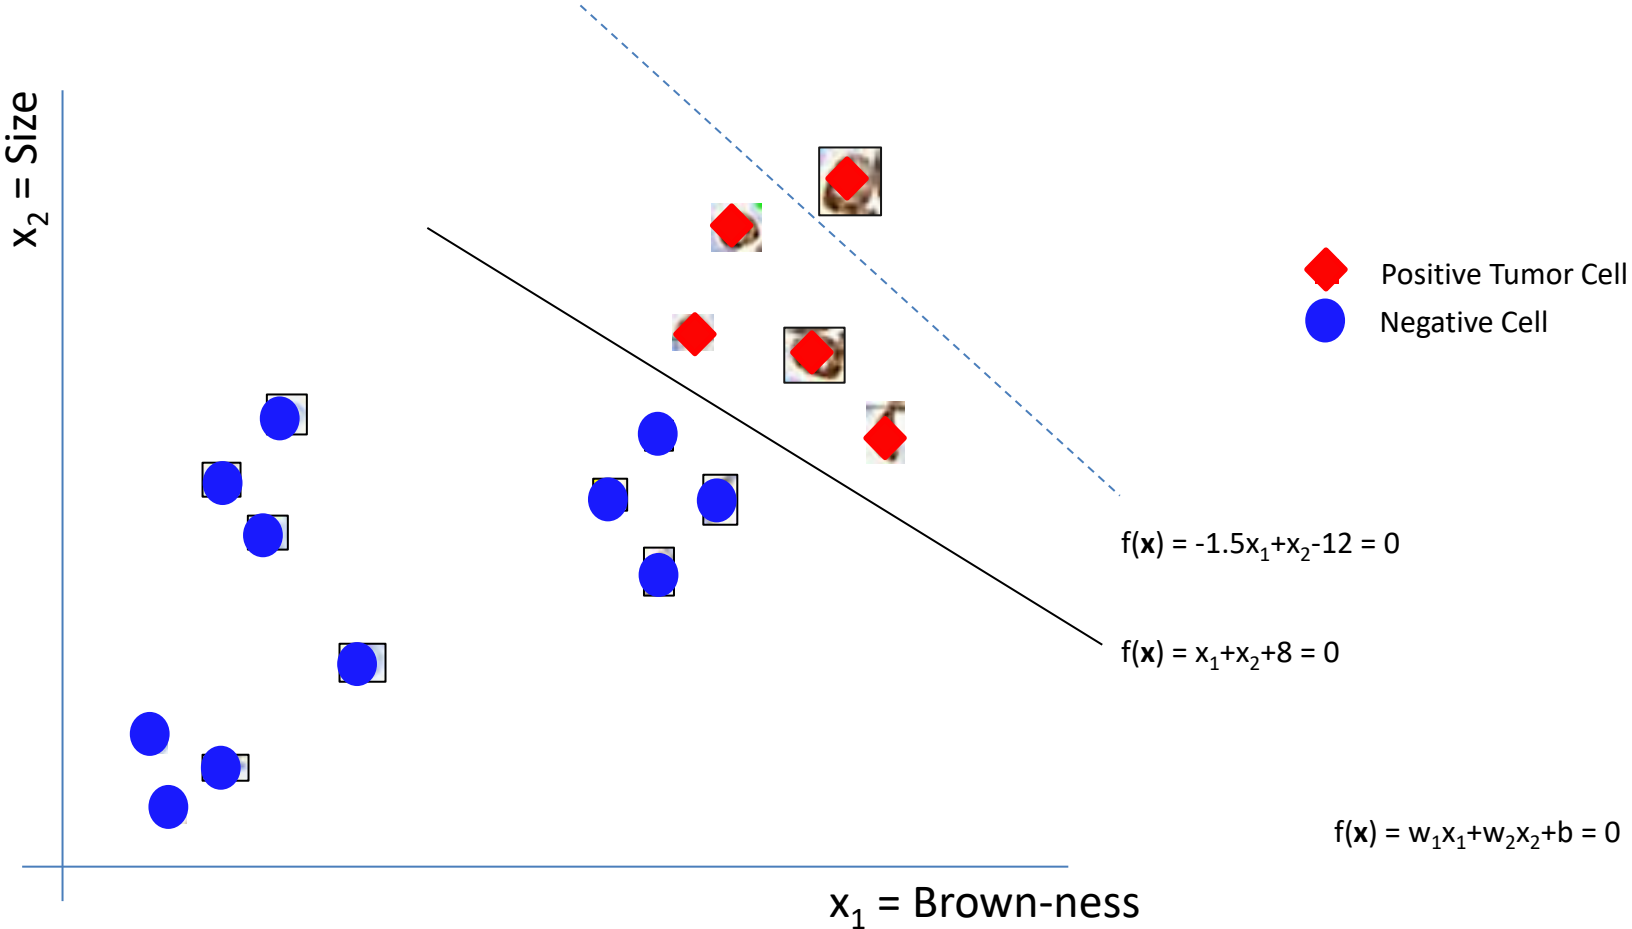


<https://info.paige.ai/prostate>



Example Independent validation study of PAIGE Prostate: Kanan, Christopher, Jillian Sue, Leo Grady, Thomas J. Fuchs, Sarat Chandarlapaty, Jorge S. Reis-Filho, Paulo G O Salles, Leonard Medeiros da Silva, Carlos Gil Ferreira, and Emilio Marcelo Pereira. "Independent Validation of Paige Prostate: Assessing Clinical Benefit of an Artificial Intelligence Tool within a Digital Diagnostic Pathology Laboratory Workflow." *Journal of Clinical Oncology* 38, no. 15_suppl (May 20, 2020): e14076–e14076. https://doi.org/10.1200/JCO.2020.38.15_suppl.e14076.

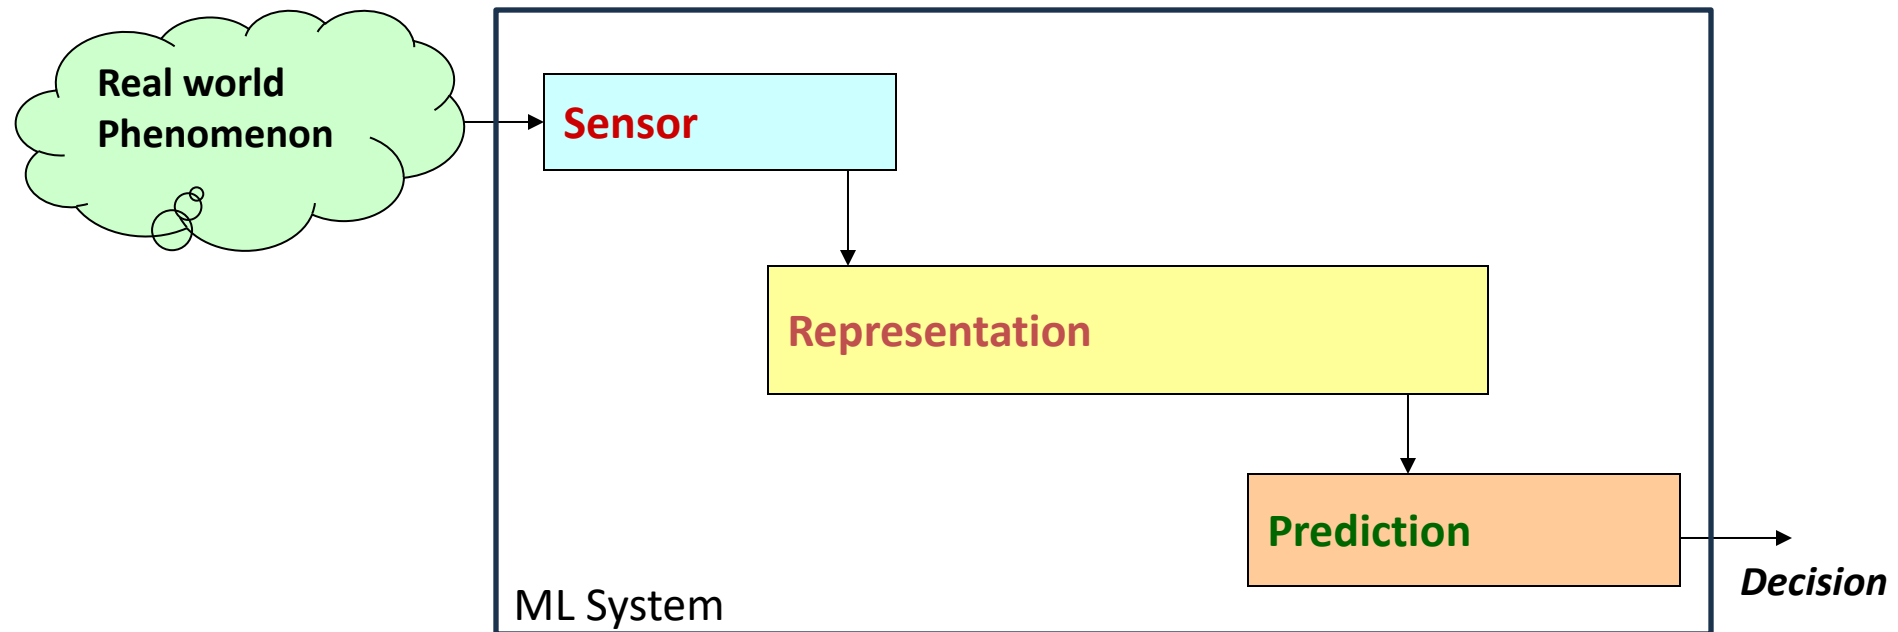
Machine Learning in Cpath with simple lines



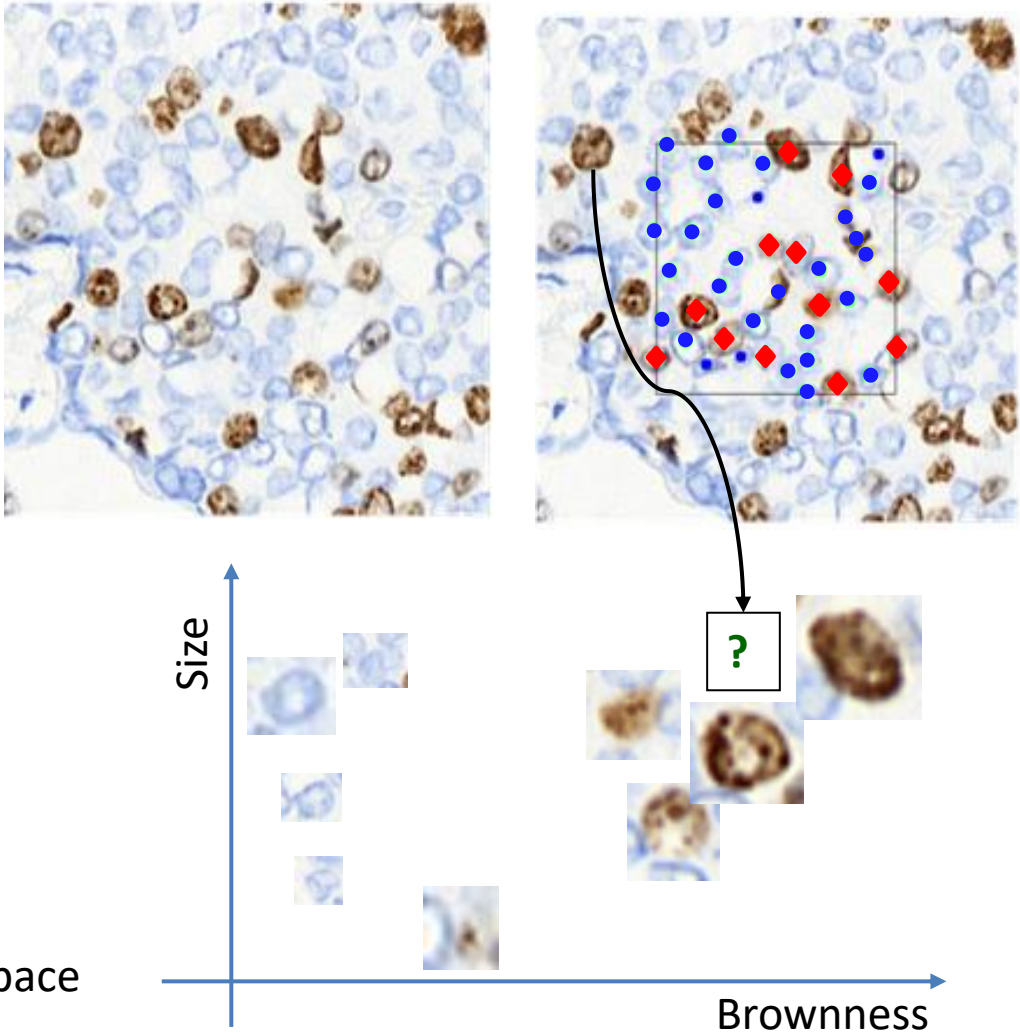

Goal is to be able to generalize to unseen data
With Deep Learning – we don't even need to define features!

Constructs of a Data Mining System for Prediction

- Identify the objective
 - Identify the unit of classification (example)
 - Image block, protein sequence,

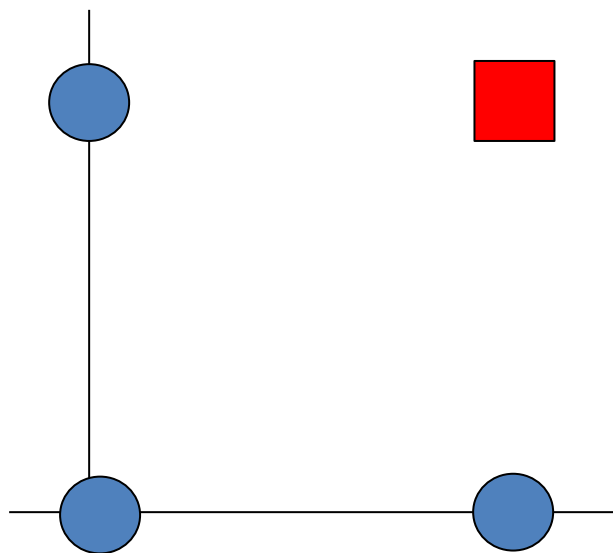


How does ML work?



Wenqi Lu, Islam M Miligy, Fayyaz Minhas, Young Saeng Park, David R J Snead, Emad A Rakha, Clare Verrill, Nasir Rajpoot "Lessons from a Breast Cell Annotation Competition Series for School Pupils." Scientific Reports, 2022. <https://ora.ox.ac.uk/objects/uuid:9e34d4e6-c677-4380-9403-759808b349aa>.

Exercise




$$f(\mathbf{x}; \mathbf{w}) = w_1 x^{(1)} + w_2 x^{(2)} + \dots + w_d x^{(d)} + b = 0$$

such that


If $y_i = +1$, then $f(\mathbf{x}_i; \mathbf{w}) > 0$

And if $y_i = -1$, then $f(\mathbf{x}_i; \mathbf{w}) < 0$

Exercise




Exercise



Doing it interactively




AND Classification Problem



<https://foxtrotmike.github.io/CS909/AND-NEURON.html>

Example (Graphical Approach to finding the discriminant function)



$$f(\mathbf{x}; \mathbf{w}) = w_1 x^{(1)} + w_2 x^{(2)} + b = 0$$

$$w_1(1.5) + w_2(0) + b = 0$$

$$b = -1.5w_1$$

$$w_1(0) + w_2(1.5) + b = 0$$

$$b = -1.5w_2$$

If I set $w_1 = 1.0$

$$b = -1.5$$


$$w_2 = 1.0$$

$$f(\mathbf{x}; \boldsymbol{\theta}) = x^{(1)} + x^{(2)} - 1.5$$

$$(1, 1) \rightarrow 0.5$$

$$(1, 0) \rightarrow -0.5, (0, 1) \rightarrow -0.5, (0, 0) \rightarrow -1.5$$

Example: Another Way (Algebraic Constraint Satisfaction to find equation)



$$f(\mathbf{x};\mathbf{w}) = w_1x^{(1)}+w_2x^{(2)}+b = 0$$

$$(1,1): w_1(1.0)+w_2(1.0)+b > 0$$

$$(1,0): w_1(1.0)+w_2(0.0)+b < 0$$

$$(0,1): w_1(0.0)+w_2(1.0)+b < 0$$

$$(0,0): w_1(0.0)+w_2(0.0)+b < 0$$

$$w_1+w_2+b > 0$$

$$w_1+b < 0$$

$$w_2+b < 0$$

$$b < 0$$


$$b = -1.5$$

$$w_1 = 1.0$$

$$w_2 = 1.0$$

Exercise

- Is this problem linearly separable?




What about this one?

- (0,0,0): -1
- (1,0,0): +1
- (0,1,0): -1
- (0,0,1): +1
- (1,0,1): +1
- (1,1,0): +1
- (0,1,1): +1
- (1,1,1): +1

$$f(\mathbf{x};\mathbf{w}) = w_1x^{(1)}+w_2x^{(2)}+w_3x^{(3)}+b = 0$$

- What about this one?



$$(1,1): w_1(1.0)+w_2(1.0)+b > 0$$

$$(1,0): w_1(1.0)+w_2(0.0)+b < 0$$


$$(0,1): w_1(0.0)+w_2(1.0)+b < 0$$

$$(0,0): w_1(0.0)+w_2(0.0)+b > 0$$

Machine learning and deep learning involve discovering meaningful representations of input data and then using these representations to partition data for various tasks

Another way of looking at Classification

- We would like to minimize the number of errors a discriminant function $f(\mathbf{x})$ makes
- **Representation:** Assume we look at only linear functions
$$f(\mathbf{x}; \mathbf{w}) = w_0 + w_1x^{(1)} + w_2x^{(2)} + \dots + w_dx^{(d)} = 0$$
- **Evaluation:** We need to define error that a particular $f(\mathbf{x}; \mathbf{w})$ makes
- **Optimization:** We need to minimize the error by tuning \mathbf{w}



Building Linear Discriminants

- Representation

- Features
- Linear Function

$$f(\mathbf{x}; \mathbf{w}) = w_1x^{(1)} + w_2x^{(2)} + \dots + w_dx^{(d)} + b = 0$$


- Evaluation

- Misclassification Error

- Optimization

- Find a line that minimizes misclassifications
- How done: Visual reckoning / Constraint Satisfaction

- Why Study Linear Models?



$$\begin{aligned}(1,1): & w_1(1.0)+w_2(1.0)+b > 0 \\(1,0): & w_1(1.0)+w_2(0.0)+b < 0 \\(0,1): & w_1(0.0)+w_2(1.0)+b < 0 \\(0,0): & w_1(0.0)+w_2(0.0)+b < 0\end{aligned}$$


A more mathematical look

- Linear Discriminants
- The linear discriminant function is given by

$$f(\mathbf{x}; \mathbf{w}) = w_1x^{(1)} + w_2x^{(2)} + \dots + w_dx^{(d)} + b = \mathbf{w}^T \mathbf{x} + b$$
$$f(\mathbf{x}'; \mathbf{w}') = w_1x^{(1)} + w_2x^{(2)} + \dots + w_dx^{(d)} + b = \mathbf{w}'^T \mathbf{x}'$$

$$\mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_d \end{bmatrix} \quad \mathbf{x} = \begin{bmatrix} x^{(1)} \\ x^{(2)} \\ \vdots \\ x^{(d)} \end{bmatrix}$$

$$\mathbf{w}' = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_d \\ b \end{bmatrix} \quad \mathbf{x}' = \begin{bmatrix} x^{(1)} \\ x^{(2)} \\ \vdots \\ x^{(d)} \\ 1 \end{bmatrix}$$



Classification Loss Function

- A misclassification is an error
 - If a training example has a label of $y = +1$, then its discriminant function score $f(\mathbf{x})$ should be _____
 - If a training example has a label of $y = -1$, then its discriminant function score $f(\mathbf{x})$ should be _____
 - Thus, we have an error whenever: _____


Classification Loss Function

- A misclassification is an error
 - If a training example has a label of $y = +1$, then its discriminant function score $f(\mathbf{x})$ should be > 0
 - If a training example has a label of $y = -1$, then its discriminant function score $f(\mathbf{x})$ should be < 0
 - Thus, we have an error whenever: $yf(\mathbf{x}) < 0$

0-1 Loss/Error

- Consider a single example:

– Our error function is: $l(f(x), y) = \begin{cases} 0 & yf(x) > 0 \\ 1 & yf(x) \leq 0 \end{cases}$






Formally called the zero-one loss function

0-1 Misclassification Error

- We want to find the parameters of the discriminant that minimize the loss for all examples in training
- Issues with 0-1 loss
 - Non Differentiable
 - Leads to poor optimization
- We need a “surrogate” or approximation of the loss
 - Should be continuous
 - Should be an over-approximation of the 0-1 loss
 - Generates at least as much error as the 0-1 loss would
 - Should be convex
 - Convex loss function leads to convex optimization problems which are easier to solve as they have a single minima

Convexity: If a line connecting two points on a curve lies on or above the curves at all times







Surrogate Classification Loss

$$l(f(\mathbf{x}), y) = \begin{cases} 0 & yf(\mathbf{x}) > 1 \\ 1 - yf(\mathbf{x}) & yf(\mathbf{x}) \leq 1 \end{cases} = \begin{cases} 0 & 1 - yf(\mathbf{x}) < 0 \\ 1 - yf(\mathbf{x}) & 1 - yf(\mathbf{x}) \geq 0 \end{cases}$$

OR

$$l(f(\mathbf{x}), y) = \max(0, 1 - yf(\mathbf{x}))$$



- **Hinge Loss Function**

- A convex over-approximation of the 0-1 loss
- Adds some “margin” of error to the classification

- Prediction label can be +1 or -1 depending upon whether $f(x) > 0$ or $f(x) < 0$
- However, we incur a loss if for positive training examples $f(x) < 1$ or for negative examples $f(x) > -1$

Optimization

$$\min_{\mathbf{w}} L(\mathbf{X}, \mathbf{Y}; \mathbf{w}) = \sum_{i=1}^N \max\{0, 1 - y_i f(\mathbf{x}_i; \mathbf{w})\}$$

- How can we solve it?
 - Take the derivative and substitute to zero
 - How else can we solve it?
 - Use gradient descent

Optimization

$$\min_{\mathbf{w}} L(\mathbf{X}, \mathbf{Y}; \mathbf{w}) = \sum_{i=1}^N l(f(\mathbf{x}_i; \mathbf{w}), y_i) = \sum_{i=1}^N \max\{0, 1 - y_i f(\mathbf{x}_i; \mathbf{w})\}$$

$$\frac{\partial L}{\partial \mathbf{w}} = \sum_{i=1}^N \frac{\partial l(f(\mathbf{x}_i; \mathbf{w}), y_i)}{\partial \mathbf{w}}$$



$$\frac{\partial}{\partial \mathbf{w}} \max\{0, 1 - y(\mathbf{w}^T \mathbf{x})\} = \begin{cases} 0 & 1 - yf(\mathbf{x}; \mathbf{w}) < 0 \\ -y\mathbf{x} & \text{else} \end{cases} = \begin{cases} -y\mathbf{x} & l(f(\mathbf{x}; \mathbf{w}), y) > 0 \\ 0 & \text{else} \end{cases}$$

What happens when $yf(x) = 1$ where the function has a “kink”?



There, we can choose to define the “sub-gradient” to be the slope of any line that lies below or on the loss function itself (see dotted lines below).

Consequently defining

$\frac{\partial L}{\partial \mathbf{w}} \Big|_{yf(x)=1} = \mathbf{0}$ should work (slope of red line).



For a simple example in which $x = 1, y = 1$



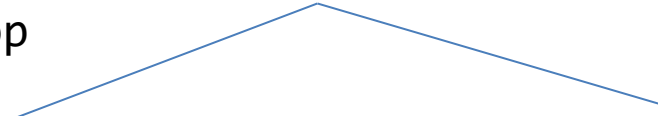
For a simple example in which $x = 1, y = 1$

Algorithm

- Given:
 - Training Examples: $\{(\mathbf{x}_i, y_i) | i = 1 \dots N\}, y_i \in \{-1, +1\}$
 - Learning rate (step size): α
- Initialize $\mathbf{w}^{(0)}$ at random
- Until Convergence ($k = 1 \dots K$ epochs)
 - For $i = 1 \dots N$
 - Pick example \mathbf{x}_i with label y_i
 - Compute $f(\mathbf{x}_i) = \mathbf{w}^{(k-1)T} \mathbf{x}_i$
 - If $y_i f(\mathbf{x}_i) < 1$ then update weight vector using gradient descent

$$\mathbf{w}^{(k)} = \mathbf{w}^{(k-1)} - \alpha \nabla l(\mathbf{w}^{(k-1)}) = \mathbf{w}^{(k-1)} - \alpha(-y_i \mathbf{x}_i) = \mathbf{w}^{(k-1)} + \alpha y_i \mathbf{x}_i$$


- Check for convergence to stop


$$\nabla_{\mathbf{w}} \max\{0, 1 - y(\mathbf{w}^T \mathbf{x})\} = \begin{cases} 0 & 1 - yf(\mathbf{x}; \mathbf{w}) < 0 \\ -y\mathbf{x} & \text{else} \end{cases}$$

REO For Perceptron

- **Representation**

- Features
- Discriminant
 - Linear



$$f(\mathbf{x}) = w_1x_1 + w_2x_2 + \dots + w_dx_d + b = \mathbf{w}^T \mathbf{x}$$

- **Evaluation**

- 0/1 (Step) Loss

$$l(f(x), y) = \begin{cases} 0 & yf(x) > 0 \\ 1 & yf(x) \leq 0 \end{cases}$$

- Hinge Loss

$$l(f(x), y) = \begin{cases} 0 & yf(x) > 1 \\ 1 - yf(x) & yf(x) \leq 1 \end{cases} = \max\{0, 1 - y(\mathbf{w}^T \mathbf{x})\}$$

- Overall Loss

$$L = \frac{1}{N} \sum_{i=1}^N l(f(\mathbf{x}_i; \mathbf{w}), y_i) \xrightarrow{iid} E[l(f(\mathbf{x}; \mathbf{w}), y)]$$

- **Optimization**

- Using Gradient Descent

$$\nabla L = \frac{1}{N} \sum_{i=1}^N \nabla_{\mathbf{w}} l(f(\mathbf{x}_i; \mathbf{w}), y_i)$$


$$\nabla_{\mathbf{w}} \max\{0, 1 - y(\mathbf{w}^T \mathbf{x})\} = \begin{cases} 0 & 1 - yf(\mathbf{x}; \mathbf{w}) < 0 \\ -y\mathbf{x} & \text{else} \end{cases} = \begin{cases} 0 & l(f(\mathbf{x}; \mathbf{w}), y) > 0 \\ -y\mathbf{x} & \text{else} \end{cases} = l(l(f(\mathbf{x}; \mathbf{w}), y))(-y\mathbf{x})$$

- Given:
 - Training Examples: $\{(\mathbf{x}_i, y_i) | i = 1 \dots N\}, y_i \in \{-1, +1\}$
- Initialize $\mathbf{w}^{(0)}$ at random
- Until Convergence ($k=1 \dots K$)
 - For $i = 1 \dots N$
 - Pick example \mathbf{x}_i with label y_i
 - Compute $f(\mathbf{x}_i) = \mathbf{w}^{(k-1)T} \mathbf{x}_i$
 - If $y_i f(\mathbf{x}_i) < 1$ then update your weight vector using gradient descent

$$\nabla_{\mathbf{w}} \max\{0, 1 - y(\mathbf{w}^T \mathbf{x})\} = \begin{cases} 0 & 1 - yf(\mathbf{x}; \mathbf{w}) < 0 \\ -y\mathbf{x} & \text{else} \end{cases}$$

$$\mathbf{w}^{(k)} = \mathbf{w}^{(k-1)} - \alpha \nabla l(\mathbf{w}^{(k-1)}) = \mathbf{w}^{(k-1)} - \alpha(-y_i \mathbf{x}_i) = \mathbf{w}^{(k-1)} + \alpha y_i \mathbf{x}_i$$

- Check for convergence to stop



$$\mathbf{w}^{(k)} \leftarrow \mathbf{w}^{(k-1)} - \alpha \nabla l(\mathbf{w}^{(k-1)})$$

$$\mathbf{w}^{(k)} \leftarrow \mathbf{w}^{(k-1)} - \alpha l(l(f(\mathbf{x}; \mathbf{w}), y))(-y\mathbf{x}) = \mathbf{w}^{(k-1)} + \alpha l(l(f(\mathbf{x}; \mathbf{w}), y))(y\mathbf{x})$$

Perceptron

- A simpler version of this algorithm is called: Perceptron
 - It updated weights whenever an example was misclassified ($y_i f(\mathbf{x}_i) < 0$) instead of when $y_i f(\mathbf{x}_i) < 1$
 - Rosenblatt (1962)
 - Minsky and Papert (1969, 1988)
 - This algorithm provides theoretical guarantees of convergence to a correct separating boundary
 - If the data is linearly separable and you allow the perceptron algorithm to run long enough, you will find the separating line!
 - **Perceptron Learning Rule Convergence Theorem**





Frank Rosenblatt
July 11, 1928 – July 11, 1971



Marvin Minsky
Aug. 9, 1927 – Jan. 24, 2016

Perceptron

- One of the first “artificial” neural networks



$$\mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_d \end{bmatrix} \quad \mathbf{x} = \begin{bmatrix} x^{(1)} \\ x^{(2)} \\ \vdots \\ x^{(d)} \end{bmatrix}$$

$$f(\mathbf{x}; \boldsymbol{\theta}) = \mathbf{w}^T \mathbf{x} + b$$

Coding Exercise

```
import numpy as np
import matplotlib.pyplot as plt
import itertools

class Perceptron:


    def __init__(self, alpha = 0.1, epochs = 200):
        self.alpha = alpha
        self.epochs = epochs
        self.W = np.array([0])
        self.bias = np.random.randn()
        self.Lambda = 0.5
    def fit(self, Xtr, Ytr):
        d = Xtr.shape[1]
        self.W = np.random.randn(d)
        for e in range(self.epochs):
            finished = True
            for i, x in enumerate(Xtr):
                if Ytr[i]*self.predict(np.atleast_2d(x))<1:
                    finished = False
                    self.W += self.alpha*Ytr[i]*x
                    self.bias += self.alpha*Ytr[i]
            if finished: break

    def score(self, x):
        return np.dot(x, self.W) + self.bias


    def predict(self, x):
        return np.sign(self.score(x))
```

```
if __name__=='__main__':
    from plotit import plotit
    Xtr = np.array([[ -1,0],[0,1],[4,4],[2,3]])
    ytr = np.array([-1,-1,+1,+1])
    clf = Perceptron()
    clf.fit(Xtr,ytr)
    z = clf.score(Xtr)
    print("Prediction Scores:",z)
    y = clf.predict(Xtr)
    print("Prediction Labels:",y)
    plotit(Xtr,ytr,clf=clf.score,conts=[0],
           extent = [-5,+5,-5,+5])
```

https://github.com/foxtrotmike/CS909/blob/master/dm_lab_2_fm.ipynb



https://github.com/foxtrotmike/CS909/blob/master/perceptron_video.py



End of Lecture

We want to make a machine that will be proud of us.

- Danny Hillis