## 10-301/601: Introduction to Machine Learning Lecture 6 - Perceptron

Hoda Heidari, Henry Chai \& Matt Gormley

2/5/24

- Announcements:
- HW2 released 1/24, due 2/5 (today!) at 11:59 PM


## Front Matter

- HW3 released on 2/5 (today!), due 2/12 at 11:59 PM
- HW3 is a written-only homework
- You may only use at most 2 late days on HW3

Q \& A:
After we do model selection using a validation dataset, should we train a final model using both the training and the validation datasets?

- Yes, absolutely! So really the sketch from last lecture should look something like:

1. Split $\mathcal{D}$ into $\mathcal{D}_{\text {train }} \cup \mathcal{D}_{\text {val }} \cup \mathcal{D}_{\text {test }}$
2. Learn classifiers using $\mathcal{D}_{\text {train }}$
3. Evaluate models using $\mathcal{D}_{v a l}$ and choose the one with lowest validation error:
4. Learn a new classifier from the best model using $\mathcal{D}_{\text {train }} \cup \mathcal{D}_{\text {val }}$
5. Optionally, use $\mathcal{D}_{\text {test }}$ to estimate the true error

- Yes! We can either:

Q \& A:
Can we use kNNs with categorical features?

1. Convert categorical features into binary ones:

2. Use a distance metric that works over categorical features e.g., the Hamming distance:

$$
d\left(\boldsymbol{x}, \boldsymbol{x}^{\prime}\right)=\sum_{d=1}^{D} \mathbb{1}\left(x_{d}=x_{d}^{\prime}\right)
$$

- See HW3 for an example of this


## Hyperparameter Optimization

- Given $\mathcal{D}=\mathcal{D}_{\text {train }} \cup \mathcal{D}_{\text {val }} \cup \mathcal{D}_{\text {test }}$, suppose we have multiple candidate hyperparameter settings:

$$
\theta_{1}, \theta_{2}, \ldots, \theta_{M}
$$

- Learn a classifier for each setting using only $\mathcal{D}_{\text {train }}$ :

$$
h_{1}, h_{2}, \ldots, h_{M}
$$

- Evaluate each one using $\mathcal{D}_{\text {val }}$ and choose the one with lowest validation error:

$$
\widehat{m}=\underset{m \in\{1, \ldots, M\}}{\operatorname{argmin}} \operatorname{err}\left(h_{m}, \mathcal{D}_{\text {val }}\right)
$$

- Now $\operatorname{err}\left(h_{\widehat{m}}, \mathcal{D}_{\text {test }}\right)$ is a good estimate of $\operatorname{err}\left(h_{\widehat{m}}\right)$ !


## How to pick hyperparameter settings to try?

- Given $\mathcal{D}=\mathcal{D}_{\text {train }} \cup \mathcal{D}_{\text {val }} \cup \mathcal{D}_{\text {test }}$, suppose we have multiple candidate hyperparameter settings:

$$
\theta_{1}, \theta_{2}, \ldots, \theta_{M}
$$

- Learn a classifier for each setting using only $\mathcal{D}_{\text {train }}$ :

$$
h_{1}, h_{2}, \ldots, h_{M}
$$

- Evaluate each one using $\mathcal{D}_{v a l}$ and choose the one with lowest validation error:

$$
\widehat{m}=\underset{m \in\{1, \ldots, M\}}{\operatorname{argmin}} \operatorname{err}\left(h_{m}, \mathcal{D}_{v a l}\right)
$$

- Now $\operatorname{err}\left(h_{\widehat{m}}, \mathcal{D}_{\text {test }}\right)$ is a good estimate of $\operatorname{err}\left(h_{\widehat{m}}\right)$ !


## General Methods for Hyperparameter Optimization

- Idea: set the hyperparameters to optimize some performance metric of the model
- Issue: if we have many hyperparameters that can all take on lots of different values, we might not be able to test all possible combinations
- Commonly used methods:
- Grid search
- Random search
- Bayesian optimization (used by Google DeepMind to optimize the hyperparameters of AlphaGo: https://arxiv.org/pdf/1812.06855v1.pdf)
- Evolutionary algorithms
- Graduate-student descent


## General Methods for Hyperparameter Optimization

- Idea: set the hyperparameters to optimize some performance metric of the model
- Issue: if we have many hyperparameters that can all take on lots of different values, we might not be able to test all possible combinations
- Commonly used methods:
- Grid search
- Random search
- Bayesian optimization (used by Google DeepMind to optimize the hyperparameters of AlphaGo: https://arxiv.org/pdf/1812.06855v1.pdf)
- Evolutionary algorithms
- Graduate-student descent


## Grid Search vs. Random Search (Bergstra and Bengio, 2012)

Grid Layout


## Poll Question 1:

Which hyperparameter optimization method do you think will perform better?

## Grid Layout

Random Layout

A. Graduate student descent (TOXIC)
B. Grid search
C. Random search

## Grid Search vs. Random Search (Bergstra and Bengio, 2012)



Important parameter

Random Layout


Important parameter

Grid and random search of nine trials for optimizing a function $f(x, y)=g(x)+h(y) \approx g(x)$ with low effective dimensionality. Above each square $g(x)$ is shown in green, and left of each square $h(y)$ is shown in yellow. With grid search, nine trials only test $g(x)$ in three distinct places. With random search, all nine trials explore distinct values of $g$. This failure of grid search is the rule rather than the exception in high dimensional hyper-parameter optimization.

# Model <br> Selection Learning Objectives 

You should be able to...

- Plan an experiment that uses training, validation, and test datasets to predict the performance of a classifier on unseen data (without cheating)
- Explain the difference between (1) training error, (2) validation error, (3) cross-validation error, (4) test error, and (5) true error
- For a given learning technique, identify the model, learning algorithm, parameters, and hyperparamters
- Select an appropriate algorithm for optimizing (aka. learning) hyperparameters

Recall:
Fisher Iris
Dataset


- Notation: in this class vectors will be assumed to be column vectors by default, i.e.,

$$
\boldsymbol{a}=\left[\begin{array}{c}
a_{1} \\
a_{2} \\
\vdots \\
a_{D}
\end{array}\right] \text { and } \boldsymbol{a}^{T}=\left[\begin{array}{llll}
a_{1} & a_{2} & \cdots & a_{D}
\end{array}\right]
$$

- The dot product between two $D$-dimensional vectors is

$$
\boldsymbol{a}^{T} \boldsymbol{b}=\left[\begin{array}{llll}
a_{1} & a_{2} & \cdots & a_{D}
\end{array}\right]\left[\begin{array}{c}
b_{1} \\
b_{2} \\
\vdots \\
b_{D}
\end{array}\right]=\sum_{d=1}^{D} a_{d} b_{d}
$$

- The L2-norm of $\boldsymbol{a}=\|\boldsymbol{a}\|_{2}=\sqrt{\boldsymbol{a}^{T} \boldsymbol{a}}$
- Two vectors are orthogonal iff

$$
\boldsymbol{a}^{T} \boldsymbol{b}=0
$$

1. On the axes below, draw the region corresponding to

$$
w_{1} x_{1}+w_{2} x_{2}+b>0
$$

where $w_{1}=1, w_{2}=2$ and $b=-4$.
2. Then draw the vector $w=\left[\begin{array}{l}w_{1} \\ w_{2}\end{array}\right]$

## Geometry Warm-up



- In 2 dimensions, $w_{1} x_{1}+w_{2} x_{2}+b=0$ defines a line
- In 3 dimensions, $w_{1} x_{1}+w_{2} x_{2}+w_{3} x_{3}+b=0$ defines a plane
- In 4+ dimensions, $\boldsymbol{w}^{T} \boldsymbol{x}+b=0$ defines a hyperplane


## Linear Decision Boundaries

- The vector $\boldsymbol{w}$ is always orthogonal to this hyperplane and always points in the direction where $\boldsymbol{w}^{T} \boldsymbol{x}+b>0$ !
- A hyperplane creates two halfspaces:
- $\mathcal{S}_{+}=\left\{\boldsymbol{x}: \boldsymbol{w}^{T} \boldsymbol{x}+b>0\right\}$ or all $\boldsymbol{x}$ s.t. $\boldsymbol{w}^{T} \boldsymbol{x}+b$ is positive
- $\mathcal{S}_{-}=\left\{\boldsymbol{x}: \boldsymbol{w}^{T} \boldsymbol{x}+b<0\right\}$ or all $\boldsymbol{x}$ s.t. $\boldsymbol{w}^{T} \boldsymbol{x}+b$ is negative


## Linear Decision Boundaries:

 Example

Goal: learn classifiers of the form $h(\boldsymbol{x})=$ $\operatorname{sign}\left(\boldsymbol{w}^{T} \boldsymbol{x}+b\right)$ (assuming $y \in\{-1,+1\})$

Key question: how do we learn the parameters, $w$ and $b$ ?

- So far, we've been learning in the batch setting, where we have access to the entire training dataset at once
- A common alternative is the online setting, where data points arrive gradually over time and we learn continuously
- Examples of online learning:
- Predicting stock prices
- Recommender systems
- Medical diagnosis
- Robotics
- For $t=1,2,3, \ldots$
- Receive an unlabeled data point, $\boldsymbol{x}^{(t)}$
- Predict its label, $\hat{y}=h_{w, b}\left(\boldsymbol{x}^{(t)}\right)$
- Observe its true label, $y^{(t)}$
- Pay a penalty if we made a mistake, $\hat{y} \neq y^{(t)}$
- Update the parameters, $\boldsymbol{w}$ and $b$
- Goal: minimize the number of mistakes made
- Initialize the weight vector and intercept to all zeros:

$$
\boldsymbol{w}=\left[\begin{array}{llll}
0 & 0 & \cdots & 0
\end{array}\right] \text { and } b=0
$$

- For $t=1,2,3, \ldots$
- Receive an unlabeled data point, $\boldsymbol{x}^{(t)}$
(Online)
Perceptron Learning Algorithm
- Predict its label, $\hat{y}=\operatorname{sign}\left(\boldsymbol{w}^{T} \boldsymbol{x}+b\right)=\left\{\begin{array}{l}+1 \text { if } \boldsymbol{w}^{T} \boldsymbol{x}+b \geq 0 \\ -1 \text { otherwise }\end{array}\right.$
- Observe its true label, $y^{(t)}$
- If we misclassified a positive point $\left(y^{(t)}=+1, \hat{y}=-1\right)$ :

$$
\begin{aligned}
& \cdot \boldsymbol{w} \leftarrow \boldsymbol{w}+\boldsymbol{x}^{(t)} \\
& \cdot b \leftarrow b+1
\end{aligned}
$$

- If we misclassified a negative point $\left(y^{(t)}=-1, \hat{y}=+1\right)$ :

$$
\begin{aligned}
& \cdot \boldsymbol{w} \leftarrow \boldsymbol{w}-\boldsymbol{x}^{(t)} \\
& \cdot b \leftarrow b-1
\end{aligned}
$$

- Initialize the weight vector and intercept to all zeros:

$$
\boldsymbol{w}=\left[\begin{array}{llll}
0 & 0 & \cdots & 0
\end{array}\right] \text { and } b=0
$$

- For $t=1,2,3, \ldots$
- Receive an unlabeled data point, $\boldsymbol{x}^{(t)}$
(Online)
Perceptron Learning Algorithm
- Predict its label, $\hat{\boldsymbol{y}}=\operatorname{sign}\left(\boldsymbol{w}^{T} \boldsymbol{x}+b\right)=\left\{\begin{array}{l}+1 \text { if } \boldsymbol{w}^{T} \boldsymbol{x}+b \geq 0 \\ -1 \text { otherwise }\end{array}\right.$
- Observe its true label, $y^{(t)}$
- If we misclassified a point $\left(y^{(t)} \neq \hat{y}\right)$ :

$$
\begin{aligned}
& \boldsymbol{w} \leftarrow \boldsymbol{w}+y^{(t)} \boldsymbol{x}^{(t)} \\
& \cdot b \leftarrow b+y^{(t)}
\end{aligned}
$$ <br> \title{

## (Online) <br> \title{ \section*{(Online) <br> <br> <br> Perceptron <br> <br> <br> Perceptron <br> <br> <br> Learning <br> <br> <br> Learning <br> <br> <br> Algorithm: <br> <br> <br> Algorithm: <br> <br> <br> Example <br> <br> <br> Example (no Intercept)} 

 (no Intercept)}}

| $x_{1}$ | $x_{2}$ | $\hat{y}$ | $y$ | Mistake? |
| :---: | :---: | :---: | :---: | :---: |
| -1 | 2 | + | - | Yes |

$x_{2}$
$\boldsymbol{w}=\left[\begin{array}{l}0 \\ 0\end{array}\right]$


# (Online) <br> Perceptron <br> Learning <br> Algorithm: <br> Example (no Intercept) 

| $x_{1}$ | $x_{2}$ | $\hat{y}$ | $y$ | Mistake? |
| :---: | :---: | :---: | :---: | :---: |
| -1 | 2 | + | - | Yes |



$$
\boldsymbol{w} \leftarrow \boldsymbol{w}+y^{(1)} \boldsymbol{x}^{(1)}=\left[\begin{array}{l}
0 \\
0
\end{array}\right]-\left[\begin{array}{c}
-1 \\
2
\end{array}\right]=\left[\begin{array}{c}
1 \\
-2
\end{array}\right]
$$

## (Online) Perceptron <br> Learning Algorithm: Example (no Intercept)

| $x_{1}$ | $x_{2}$ | $\hat{y}$ | $y$ | Mistake? |
| :---: | :---: | :---: | :---: | :---: |
| -1 | 2 | + | - | Yes |
| 1 | 0 | + | + | No |



## (Online) <br> Perceptron <br> Learning <br> Algorithm: <br> Example (no Intercept)

| $x_{1}$ | $x_{2}$ | $\widehat{y}$ | $y$ | Mistake? |
| :---: | :---: | :---: | :---: | :---: |
| -1 | 2 | + | - | Yes |
| 1 | 0 | + | + | No |
| 1 | 1 | - | + | Yes |

$$
\begin{aligned}
& \boldsymbol{w}=\left[\begin{array}{c}
1 \\
-2
\end{array}\right] \\
& \boldsymbol{w} \leftarrow \boldsymbol{w}+y^{(3)} \boldsymbol{x}^{(3)}=\left[\begin{array}{c}
1 \\
-2
\end{array}\right]+\left[\begin{array}{l}
1 \\
1
\end{array}\right]=\left[\begin{array}{c}
2 \\
-1
\end{array}\right]
\end{aligned}
$$



| $x_{1}$ | $x_{2}$ | $\widehat{y}$ | $y$ | Mistake? |
| :---: | :---: | :---: | :---: | :---: |
| -1 | 2 | + | - | Yes |
| 1 | 0 | + | + | No |
| 1 | 1 | - | + | Yes | <br> \section*{(Online) <br> \section*{(Online) <br> <br> Perceptron <br> <br> Perceptron <br> <br> Learning <br> <br> Learning <br> <br> Algorithm: <br> <br> Algorithm: <br> <br> Example <br> <br> Example <br> <br> (no Intercept)} <br> <br> (no Intercept)}

$$
\begin{aligned}
& \boldsymbol{w}=\left[\begin{array}{c}
1 \\
-2
\end{array}\right] \\
& \boldsymbol{w} \leftarrow \boldsymbol{w}+y^{(3)} \boldsymbol{x}^{(3)}=\left[\begin{array}{c}
1 \\
-2
\end{array}\right]+\left[\begin{array}{l}
1 \\
1
\end{array}\right]=\left[\begin{array}{c}
2 \\
-1
\end{array}\right]
\end{aligned}
$$

## (Online) <br> Perceptron <br> Learning <br> Algorithm: <br> Example (no Intercept) <br> $$
\boldsymbol{w}=\left[\begin{array}{c} 2 \\ -1 \end{array}\right]
$$

| $x_{1}$ | $x_{2}$ | $\widehat{y}$ | $y$ | Mistake? |
| :---: | :---: | :---: | :---: | :---: |
| -1 | 2 | + | - | Yes |
| 1 | 0 | + | + | No |
| 1 | 1 | - | + | Yes |
| -1 | 0 | - | - | No |


(Online)
Perceptron
Learning
Algorithm:
Example (no Intercept)

| $x_{1}$ | $x_{2}$ | $\widehat{y}$ | $y$ | Mistake? |
| :---: | :---: | :---: | :---: | :---: |
| -1 | 2 | + | - | Yes |
| 1 | 0 | + | + | No |
| 1 | 1 | - | + | Yes |
| -1 | 0 | - | - | No |
| -1 | -2 | + | - | Yes |



$$
w=\left[\begin{array}{c}
2 \\
-1
\end{array}\right]
$$

$$
\boldsymbol{w} \leftarrow \boldsymbol{w}+\boldsymbol{y}^{(5)} \boldsymbol{x}^{(5)}=\left[\begin{array}{c}
2 \\
-1
\end{array}\right]-\left[\begin{array}{c}
-1 \\
-2
\end{array}\right]=\left[\begin{array}{l}
3 \\
1
\end{array}\right]
$$

(Online)
Perceptron
Learning
Algorithm:
Example (no Intercept)

| $x_{1}$ | $x_{2}$ | $\hat{y}$ | $y$ | Mistake? |
| :---: | :---: | :---: | :---: | :---: |
| -1 | 2 | + | - | Yes |
| 1 | 0 | + | + | No |
| 1 | 1 | - | + | Yes |
| -1 | 0 | - | - | No |
| -1 | -2 | + | - | Yes |



$$
w=\left[\begin{array}{c}
2 \\
-1
\end{array}\right]
$$

$$
\boldsymbol{w} \leftarrow \boldsymbol{w}+y^{(5)} \boldsymbol{x}^{(5)}=\left[\begin{array}{c}
2 \\
-1
\end{array}\right]-\left[\begin{array}{l}
-1 \\
-2
\end{array}\right]=\left[\begin{array}{l}
3 \\
1
\end{array}\right]
$$

# (Online) Perceptron Learning Algorithm: Example (no Intercept) 

| $x_{1}$ | $x_{2}$ | $\hat{y}$ | $y$ | Mistake? |
| :---: | :---: | :---: | :---: | :---: |
| -1 | 2 | + | - | Yes |
| 1 | 0 | + | + | No |
| 1 | 1 | - | + | Yes |
| -1 | 0 | - | - | No |
| -1 | -2 | + | - | Yes |
| 1 | -1 | + | + | No |

$$
w=\left[\begin{array}{l}
3 \\
1
\end{array}\right]
$$



- The intercept shifts the decision boundary off the origin
- Increasing $b$ shifts the decision boundary towards the negative side
- Decreasing $b$ shifts the decision boundary towards the positive side



## Poll Question 2:

- True or False: Unlike Decision Trees and $k$-Nearest Neighbors, the Perceptron learning algorithm does not suffer from overfitting because it does not have any hyperparameters that could be over-tuned on the training validation data.
A. True
B. True and False (TOXIC)
C. False
- If we add a 1 to the beginning of every feature vector e.g.,

$$
\boldsymbol{x}^{\prime}=\left[\begin{array}{l}
1 \\
x
\end{array}\right]=\left[\begin{array}{c}
1 \\
x_{1} \\
x_{2} \\
\vdots \\
x_{D}
\end{array}\right] \ldots
$$

- ... we can just fold the intercept into the weight vector!

$$
\boldsymbol{\theta}=\left[\begin{array}{c}
b \\
w_{1} \\
w_{2} \\
\vdots \\
w_{D}
\end{array}\right] \rightarrow \boldsymbol{\theta}^{T} \boldsymbol{x}^{\prime}=\boldsymbol{w}^{T} \boldsymbol{x}+b
$$

- Initialize the weight vector and intercept to all zeros:

$$
\boldsymbol{w}=\left[\begin{array}{llll}
0 & 0 & \cdots & 0
\end{array}\right] \text { and } b=0
$$

- For $t=1,2,3, \ldots$
- Receive an unlabeled data point, $\boldsymbol{x}^{(t)}$
(Online)
Perceptron Learning Algorithm
- Predict its label, $\hat{y}=\operatorname{sign}\left(\boldsymbol{w}^{T} \boldsymbol{x}+b\right)=\left\{\begin{array}{l}+1 \text { if } \boldsymbol{w}^{T} \boldsymbol{x}+b \geq 0 \\ -1 \text { otherwise }\end{array}\right.$
- Observe its true label, $y^{(t)}$
- If we misclassified a point $\left(y^{(t)} \neq \hat{y}\right)$ :

$$
\begin{aligned}
& \boldsymbol{w} \leftarrow \boldsymbol{w}+y^{(t)} \boldsymbol{x}^{(t)} \\
& \cdot b \leftarrow b+y^{(t)}
\end{aligned}
$$

- Initialize the parameters to all zeros:

$$
\boldsymbol{\theta}=\left[\begin{array}{llll}
0 & 0 & \cdots & 0
\end{array}\right]
$$

- For $t=1,2,3, \ldots$
- Receive an unlabeled data point, $\boldsymbol{x}^{(t)}$


## (Online) Perceptron Learning Algorithm

- Predict its label, $\hat{y}=\operatorname{sign}\left(\boldsymbol{\theta}^{T} \boldsymbol{x}^{\prime(t)}\right)=\left\{\begin{array}{l}+1 \text { if } \boldsymbol{\theta}^{T} \boldsymbol{x}^{\prime(t)} \geq 0 \\ -1 \text { otherwise }\end{array}\right.$
- Observe its true label, $y^{(t)}$
- If we misclassified a point $\left(y^{(t)} \neq \hat{y}\right)$ :

$$
\cdot \boldsymbol{\theta} \leftarrow \boldsymbol{\theta}+y^{(t)} \boldsymbol{x}^{\prime(t)}
$$

Automatically handles updating the intercept
(Online) Perceptron
Learning Algorithm: Inductive Bias

- The decision boundary is linear and recent mistakes are more important than older ones (and should be corrected immediately)
- Initialize the parameters to all zeros:

$$
\boldsymbol{\theta}=\left[\begin{array}{llll}
0 & 0 & \cdots & 0
\end{array}\right]
$$

(Online) Perceptron Learning Algorithm

- For $t=1,2,3, \ldots$
- Receive an unlabeled data point, $\boldsymbol{x}^{(t)}$
- Predict its label, $\hat{y}=\operatorname{sign}\left(\boldsymbol{\theta}^{T} \boldsymbol{x}^{\prime(t)}\right)$
- Observe its true label, $y^{(t)}$
- If we misclassified a point $\left(y^{(t)} \neq \hat{y}\right)$ :

$$
\text { - } \boldsymbol{\theta} \leftarrow \boldsymbol{\theta}+y^{(t)} \boldsymbol{x}^{\prime(t)}
$$

- Input: $\mathcal{D}=\left\{\left(\boldsymbol{x}^{(1)}, y^{(1)}\right),\left(\boldsymbol{x}^{(2)}, y^{(2)}\right), \ldots,\left(\boldsymbol{x}^{(N)}, y^{(N)}\right)\right\}$
- Initialize the parameters to all zeros:

$$
\boldsymbol{\theta}=\left[\begin{array}{llll}
0 & 0 & \cdots & 0
\end{array}\right]
$$

(Batch)
Perceptron Learning Algorithm

- While NOT CONVERGED
- For $t \in\{1, \ldots, N\}$
- Predict the label of $\boldsymbol{x}^{\prime(t)}, \hat{y}=\operatorname{sign}\left(\boldsymbol{\theta}^{T} \boldsymbol{x}^{\prime(t)}\right)$
- Observe its true label, $y^{(t)}$
- If we misclassified $\boldsymbol{x}^{\prime(t)}\left(y^{(t)} \neq \hat{y}\right)$ :

$$
\cdot \boldsymbol{\theta} \leftarrow \boldsymbol{\theta}+y^{(t)} \boldsymbol{x}^{\prime(t)}
$$

- True or False: The parameter vector $\boldsymbol{w}$ learned by the batch Perceptron Learning Algorithm can be written as a linear combination of the examples, i.e.,

$$
\boldsymbol{w}=c_{1} \boldsymbol{x}^{(1)}+c_{2} \boldsymbol{x}^{(2)}+\cdots+c_{N} \boldsymbol{x}^{(M)}
$$

A. True and False (TOXIC)
B. True
C. False

- Definitions:
- A dataset $\mathcal{D}$ is linearly separable if $\exists$ a linear decision boundary that perfectly classifies the data points in $\mathcal{D}$
- The margin, $\gamma$, of a dataset $\mathcal{D}$ is the greatest possible distance between a linear separator and the closest data point in $\mathcal{D}$ to that linear separator

- Theorem: if the data points seen by the Perceptron Learning Algorithm (online and batch)

1. lie in a ball of radius $R$ (centered around the origin)
2. have a margin of $\gamma$
then the algorithm makes at most $(R / \gamma)^{2}$ mistakes.

- Key Takeaway: if the training dataset is linearly separable, the batch Perceptron Learning Algorithm will converge (i.e., stop making mistakes on the training dataset or achieve 0 training error) in a finite number of steps!


## Computing the Margin

- Let $\boldsymbol{x}^{\prime}$ be an arbitrary point on the hyperplane $\boldsymbol{w}^{T} \boldsymbol{x}+b=0$ and let $\boldsymbol{x}$ " be an arbitrary point
- The distance between $\boldsymbol{x}$ " and $\boldsymbol{w}^{T} \boldsymbol{x}+b=0$ is equal to the magnitude of the projection of $x^{\prime \prime}-x^{\prime}$ onto $\frac{w}{\|w\|_{2}}$, the unit vector orthogonal to the hyperplane



## Computing the Margin

- Let $\boldsymbol{x}^{\prime}$ be an arbitrary point on the hyperplane $\boldsymbol{w}^{T} \boldsymbol{x}+b=0$ and let $\boldsymbol{x}$ " be an arbitrary point
- The distance between $\boldsymbol{x}$ " and $\boldsymbol{w}^{T} \boldsymbol{x}+b=0$ is equal to the magnitude of the projection of $x^{\prime \prime}-x^{\prime}$ onto $\frac{w}{\|w\|_{2}}$, the unit vector orthogonal to the hyperplane



## Computing the Margin

- Let $\boldsymbol{x}^{\prime}$ be an arbitrary point on the hyperplane $\boldsymbol{w}^{T} \boldsymbol{x}+b=0$ and let $\boldsymbol{x}$ " be an arbitrary point
- The distance between $\boldsymbol{x}$ " and $\boldsymbol{w}^{T} \boldsymbol{x}+b=0$ is equal to the magnitude of the projection of $x^{\prime \prime}-x^{\prime}$ onto $\frac{w}{\|w\|_{2}}$, the unit vector orthogonal to the hyperplane

- Let $\boldsymbol{x}^{\prime}$ be an arbitrary point on the hyperplane and let $x^{\prime \prime}$ be an arbitrary point
- The distance between $\boldsymbol{x}$ " and $\boldsymbol{w}^{T} \boldsymbol{x}+b=0$ is equal to the magnitude of the projection of $x^{\prime \prime}-x^{\prime}$ onto $\frac{w}{\|w\|_{2}}$, the unit vector orthogonal to the hyperplane

$$
\left|\frac{\boldsymbol{w}^{T}\left(\boldsymbol{x}^{\prime \prime}-\boldsymbol{x}^{\prime}\right)}{\|\boldsymbol{w}\|_{2}}\right|=\frac{\left|\boldsymbol{w}^{T} \boldsymbol{x}^{\prime \prime}-\boldsymbol{w}^{T} \boldsymbol{x}^{\prime}\right|}{\|\boldsymbol{w}\|_{2}}=\frac{\left|\boldsymbol{w}^{T} \boldsymbol{x}^{\prime \prime}+b\right|}{\|\boldsymbol{w}\|_{2}}
$$

## Computing the Margin

You should be able to...

- Explain the difference between online learning and batch learning
- Implement the perceptron algorithm for binary classification [CIML]


## Perceptron Learning Objectives

- Determine whether the perceptron algorithm will converge based on properties of the dataset, and the limitations of the convergence guarantees
- Describe the inductive bias of perceptron and the limitations of linear models
- Draw the decision boundary of a linear model
- Identify whether a dataset is linearly separable or not
- Defend the use of a bias term in perceptron (shifting points after projection onto weight vector)

