

CSCI567 Machine Learning (Fall 2020)

Prof. Haipeng Luo

U of Southern California

Sep 10, 2020

Administration

- HW 1 is due on Tue, 9/15.

Administration

- HW 1 is due on Tue, 9/15.
- Last week to enroll.

Outline

- 1 Review of Last Lecture
- 2 Linear Classifier and Surrogate Losses
- 3 Perceptron
- 4 Logistic regression

Outline

- 1 Review of Last Lecture
- 2 Linear Classifier and Surrogate Losses
- 3 Perceptron
- 4 Logistic regression

Regression

Predicting a continuous outcome variable using past observations

- temperature, amount of rainfall, house price, etc.

Key difference from classification

- continuous vs discrete
- measure *prediction errors* differently.
- lead to quite different learning algorithms.

Linear Regression: regression with linear models: $f(\mathbf{w}) = \mathbf{w}^T \mathbf{x}$

Least square solution

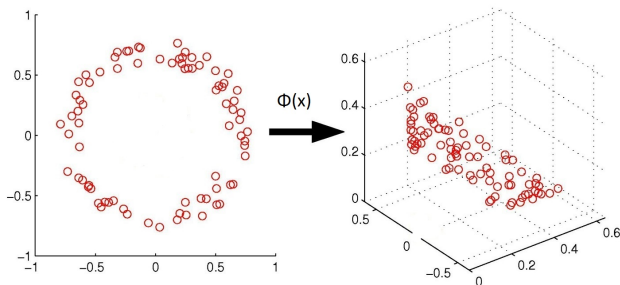
$$\begin{aligned}
 \mathbf{w}^* &= \underset{\mathbf{w}}{\operatorname{argmin}} \operatorname{RSS}(\mathbf{w}) \\
 &= \underset{\mathbf{w}}{\operatorname{argmin}} \|\mathbf{X}\mathbf{w} - \mathbf{y}\|_2^2 \\
 &= (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}
 \end{aligned}$$

$$\mathbf{X} = \begin{pmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \vdots \\ \mathbf{x}_N^T \end{pmatrix}, \quad \mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{pmatrix}$$

Two approaches to find the minimum:

- find **stationary points** by setting gradient = 0
- “**complete the square**”

Regression with nonlinear basis



Model: $f(x) = w^T \phi(x)$ where $w \in \mathbb{R}^M$

Similar least square solution: $w^* = (\Phi^T \Phi)^{-1} \Phi^T y$

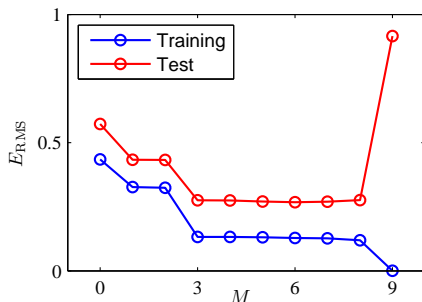
Underfitting and Overfitting

$M \leq 2$ is *underfitting* the data

- large training error
- large test error

$M \geq 9$ is *overfitting* the data

- small training error
- **large test error**



How to prevent overfitting? more data + regularization

$$\mathbf{w}^* = \underset{\mathbf{w}}{\operatorname{argmin}} \left(\operatorname{RSS}(\mathbf{w}) + \lambda \|\mathbf{w}\|_2^2 \right) = (\Phi^T \Phi + \lambda I)^{-1} \Phi^T \mathbf{y}$$

General idea to derive ML algorithms

Step 1. Pick a set of **models** \mathcal{F}

- e.g. $\mathcal{F} = \{f(\mathbf{x}) = \mathbf{w}^T \mathbf{x} \mid \mathbf{w} \in \mathbb{R}^D\}$
- e.g. $\mathcal{F} = \{f(\mathbf{x}) = \mathbf{w}^T \Phi(\mathbf{x}) \mid \mathbf{w} \in \mathbb{R}^M\}$

Step 2. Define **error/loss** $L(y', y)$

Step 3. Find **empirical risk minimizer (ERM)**:

$$\mathbf{f}^* = \operatorname{argmin}_{f \in \mathcal{F}} \sum_{n=1}^N L(f(x_n), y_n)$$

or **regularized empirical risk minimizer**:

$$\mathbf{f}^* = \operatorname{argmin}_{f \in \mathcal{F}} \sum_{n=1}^N L(f(x_n), y_n) + \lambda R(f)$$

ML becomes optimization

Outline

- 1 Review of Last Lecture
- 2 Linear Classifier and Surrogate Losses
- 3 Perceptron
- 4 Logistic regression

Classification

Recall the setup:

- input (feature vector): $\mathbf{x} \in \mathbb{R}^D$
- output (label): $y \in [C] = \{1, 2, \dots, C\}$
- goal: learn a mapping $f : \mathbb{R}^D \rightarrow [C]$

Classification

Recall the setup:

- input (feature vector): $\mathbf{x} \in \mathbb{R}^D$
- output (label): $y \in [C] = \{1, 2, \dots, C\}$
- goal: learn a mapping $f : \mathbb{R}^D \rightarrow [C]$

This lecture: **binary classification**

- Number of classes: $C = 2$
- Labels: $\{-1, +1\}$ (cat or dog, fraud or not, price up or down...)

Classification

Recall the setup:

- input (feature vector): $\mathbf{x} \in \mathbb{R}^D$
- output (label): $y \in [C] = \{1, 2, \dots, C\}$
- goal: learn a mapping $f : \mathbb{R}^D \rightarrow [C]$

This lecture: **binary classification**

- Number of classes: $C = 2$
- Labels: $\{-1, +1\}$ (cat or dog, fraud or not, price up or down...)

We have discussed **nearest neighbor classifier**:

- require carrying the training set
- more like a heuristic

Deriving classification algorithms

Let's follow the recipe:

Step 1. Pick a set of models \mathcal{F} .

Deriving classification algorithms

Let's follow the recipe:

Step 1. Pick a set of models \mathcal{F} .

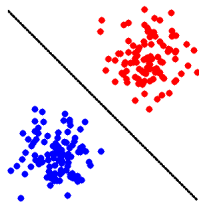
Again try linear models, but how to predict a label using $w^T x$?

Deriving classification algorithms

Let's follow the recipe:

Step 1. Pick a set of models \mathcal{F} .

Again try linear models, but how to predict a label using $w^T x$?



Deriving classification algorithms

Let's follow the recipe:

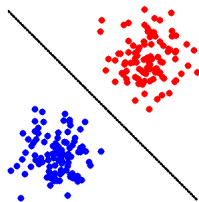
Step 1. Pick a set of models \mathcal{F} .

Again try linear models, but how to predict a label using $w^T x$?

Sign of $w^T x$ predicts the label:

$$\text{sign}(w^T x) = \begin{cases} +1 & \text{if } w^T x > 0 \\ -1 & \text{if } w^T x \leq 0 \end{cases}$$

(Sometimes use sgn for sign too.)



The models

The set of **(separating) hyperplanes**:

$$\mathcal{F} = \{f(\boldsymbol{x}) = \text{sgn}(\boldsymbol{w}^T \boldsymbol{x}) \mid \boldsymbol{w} \in \mathbb{R}^D\}$$

The models

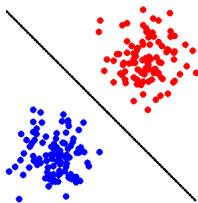
The set of **(separating) hyperplanes**:

$$\mathcal{F} = \{f(x) = \text{sgn}(w^T x) \mid w \in \mathbb{R}^D\}$$

Good choice for *linearly separable* data, i.e., $\exists w$ s.t.

$$\text{sgn}(w^T x_n) = y_n$$

for all $n \in [N]$.



The models

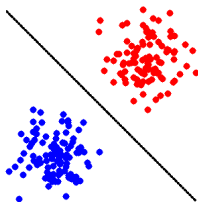
The set of **(separating) hyperplanes**:

$$\mathcal{F} = \{f(x) = \text{sgn}(\mathbf{w}^T \mathbf{x}) \mid \mathbf{w} \in \mathbb{R}^D\}$$

Good choice for *linearly separable* data, i.e., $\exists \mathbf{w}$ s.t.

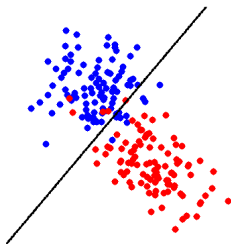
$$\text{sgn}(\mathbf{w}^T \mathbf{x}_n) = y_n \quad \text{or} \quad y_n \mathbf{w}^T \mathbf{x}_n > 0$$

for all $n \in [N]$.



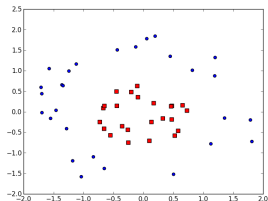
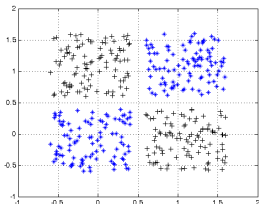
The models

Still makes sense for “almost” linearly separable data



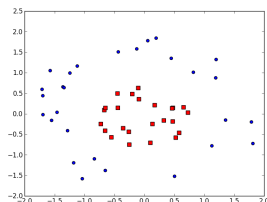
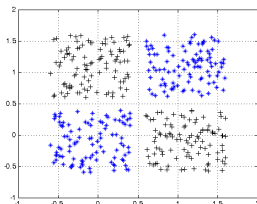
The models

For clearly not linearly separable data,



The models

For clearly not linearly separable data,



Again can apply a **nonlinear mapping** Φ :

$$\mathcal{F} = \{f(x) = \text{sgn}(w^T \Phi(x)) \mid w \in \mathbb{R}^M\}$$

More discussions in the next two lectures.

0-1 Loss

Step 2. Define error/loss $L(y', y)$.

0-1 Loss

Step 2. Define error/loss $L(y', y)$.

Most natural one for classification: **0-1 loss** $L(y', y) = \mathbb{I}[y' \neq y]$

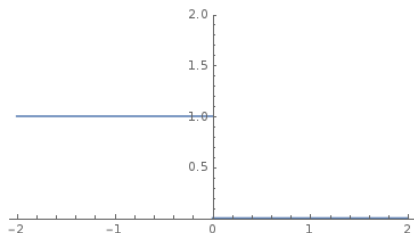
0-1 Loss

Step 2. Define error/loss $L(y', y)$.

Most natural one for classification: **0-1 loss** $L(y', y) = \mathbb{I}[y' \neq y]$

For classification, more convenient to look at the loss **as a function of** $yw^T x$. That is, with

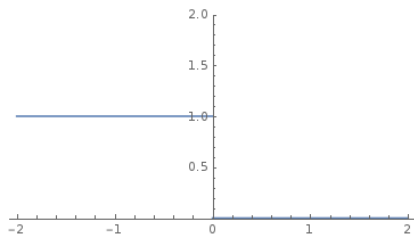
$$\ell_{0-1}(z) = \mathbb{I}[z \leq 0]$$



the loss for hyperplane w on example (x, y) is $\ell_{0-1}(yw^T x)$

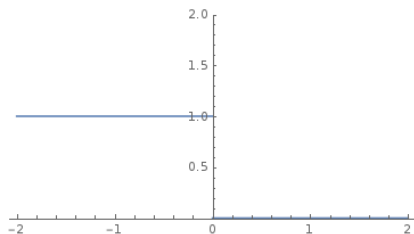
Minimizing 0-1 loss is hard

However, 0-1 loss is *not convex*.



Minimizing 0-1 loss is hard

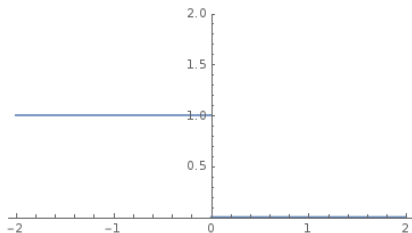
However, 0-1 loss is *not convex*.



Even worse, minimizing 0-1 loss is *NP-hard in general*.

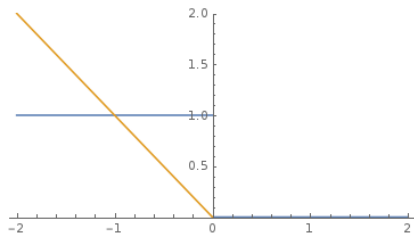
Surrogate Losses

Solution: find a **convex surrogate loss**



Surrogate Losses

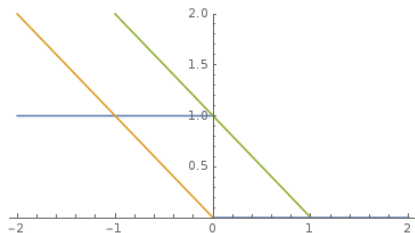
Solution: find a **convex surrogate loss**



- **perceptron loss** $\ell_{\text{perceptron}}(z) = \max\{0, -z\}$ (used in Perceptron)

Surrogate Losses

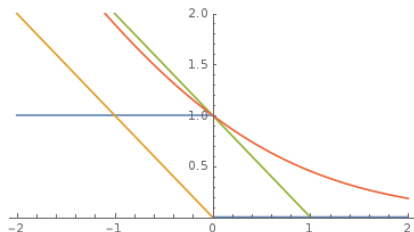
Solution: find a **convex surrogate loss**



- **perceptron loss** $\ell_{\text{perceptron}}(z) = \max\{0, -z\}$ (used in Perceptron)
- **hinge loss** $\ell_{\text{hinge}}(z) = \max\{0, 1 - z\}$ (used in SVM and many others)

Surrogate Losses

Solution: find a **convex surrogate loss**



- **perceptron loss** $\ell_{\text{perceptron}}(z) = \max\{0, -z\}$ (used in Perceptron)
- **hinge loss** $\ell_{\text{hinge}}(z) = \max\{0, 1 - z\}$ (used in SVM and many others)
- **logistic loss** $\ell_{\text{logistic}}(z) = \log(1 + \exp(-z))$ (used in logistic regression; the base of \log doesn't matter)

ML becomes convex optimization

Step 3. Find ERM:

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^D} \sum_{n=1}^N \ell(y_n \mathbf{w}^T \mathbf{x}_n) = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^D} \frac{1}{N} \sum_{n=1}^N \ell(y_n \mathbf{w}^T \mathbf{x}_n)$$

where $\ell(\cdot)$ can be perceptron/hinge/logistic loss

ML becomes convex optimization

Step 3. Find ERM:

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^D} \sum_{n=1}^N \ell(y_n \mathbf{w}^T \mathbf{x}_n) = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^D} \frac{1}{N} \sum_{n=1}^N \ell(y_n \mathbf{w}^T \mathbf{x}_n)$$

where $\ell(\cdot)$ can be perceptron/hinge/logistic loss

- *no closed-form* in general (unlike linear regression)

ML becomes convex optimization

Step 3. Find ERM:

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^D} \sum_{n=1}^N \ell(y_n \mathbf{w}^T \mathbf{x}_n) = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^D} \frac{1}{N} \sum_{n=1}^N \ell(y_n \mathbf{w}^T \mathbf{x}_n)$$

where $\ell(\cdot)$ can be perceptron/hinge/logistic loss

- *no closed-form* in general (unlike linear regression)
- can apply general convex optimization methods

ML becomes convex optimization

Step 3. Find ERM:

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^D} \sum_{n=1}^N \ell(y_n \mathbf{w}^T \mathbf{x}_n) = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^D} \frac{1}{N} \sum_{n=1}^N \ell(y_n \mathbf{w}^T \mathbf{x}_n)$$

where $\ell(\cdot)$ can be perceptron/hinge/logistic loss

- *no closed-form* in general (unlike linear regression)
- can apply general convex optimization methods

Note: minimizing perceptron loss *does not really make sense* (try $\mathbf{w} = \mathbf{0}$), but the algorithm derived from this perspective does.

Outline

- 1 Review of Last Lecture
- 2 Linear Classifier and Surrogate Losses
- 3 **Perceptron**
 - Numerical optimization
 - Applying (S)GD to perceptron loss
- 4 Logistic regression

The Perceptron Algorithm

In one sentence: Stochastic Gradient Descent applied to perceptron loss

The Perceptron Algorithm

In one sentence: Stochastic Gradient Descent applied to perceptron loss

i.e. find the minimizer of

$$\begin{aligned} F(\mathbf{w}) &= \frac{1}{N} \sum_{n=1}^N \ell_{\text{perceptron}}(y_n \mathbf{w}^T \mathbf{x}_n) \\ &= \frac{1}{N} \sum_{n=1}^N \max\{0, -y_n \mathbf{w}^T \mathbf{x}_n\} \end{aligned}$$

using SGD

A detour of numerical optimization methods

We describe two simple yet extremely popular methods

- **Gradient Descent (GD)**: simple and fundamental
- **Stochastic Gradient Descent (SGD)**: faster, effective for large-scale problems

A detour of numerical optimization methods

We describe two simple yet extremely popular methods

- **Gradient Descent (GD)**: simple and fundamental
- **Stochastic Gradient Descent (SGD)**: faster, effective for large-scale problems

Gradient is sometimes referred to as *first-order* information of a function. Therefore, these methods are called *first-order methods*.

Gradient Descent (GD)

Goal: minimize $F(\boldsymbol{w})$

Gradient Descent (GD)

Goal: minimize $F(\boldsymbol{w})$

Algorithm: keep moving in the *negative gradient direction*

Gradient Descent (GD)

Goal: minimize $F(\mathbf{w})$

Algorithm: keep moving in the *negative gradient direction*

Start from some $\mathbf{w}^{(0)}$. For $t = 0, 1, 2, \dots$

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \eta \nabla F(\mathbf{w}^{(t)})$$

where $\eta > 0$ is called step size or learning rate

Gradient Descent (GD)

Goal: minimize $F(\mathbf{w})$

Algorithm: keep moving in the *negative gradient direction*

Start from some $\mathbf{w}^{(0)}$. For $t = 0, 1, 2, \dots$

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \eta \nabla F(\mathbf{w}^{(t)})$$

where $\eta > 0$ is called step size or learning rate

- in theory η should be set in terms of some parameters of F
- in practice we just try several small values

An example

Example: $F(\mathbf{w}) = 0.5(w_1^2 - w_2)^2 + 0.5(w_1 - 1)^2$.

An example

Example: $F(\mathbf{w}) = 0.5(w_1^2 - w_2)^2 + 0.5(w_1 - 1)^2$. Gradient is

$$\frac{\partial F}{\partial w_1} = 2(w_1^2 - w_2)w_1 + w_1 - 1 \qquad \frac{\partial F}{\partial w_2} = -(w_1^2 - w_2)$$

An example

Example: $F(\mathbf{w}) = 0.5(w_1^2 - w_2)^2 + 0.5(w_1 - 1)^2$. Gradient is

$$\frac{\partial F}{\partial w_1} = 2(w_1^2 - w_2)w_1 + w_1 - 1 \qquad \frac{\partial F}{\partial w_2} = -(w_1^2 - w_2)$$

GD:

- Initialize $w_1^{(0)}$ and $w_2^{(0)}$ (to be 0 or *randomly*), $t = 0$

An example

Example: $F(\mathbf{w}) = 0.5(w_1^2 - w_2)^2 + 0.5(w_1 - 1)^2$. Gradient is

$$\frac{\partial F}{\partial w_1} = 2(w_1^2 - w_2)w_1 + w_1 - 1 \quad \frac{\partial F}{\partial w_2} = -(w_1^2 - w_2)$$

GD:

- Initialize $w_1^{(0)}$ and $w_2^{(0)}$ (to be 0 or *randomly*), $t = 0$
- do

$$w_1^{(t+1)} \leftarrow w_1^{(t)} - \eta \left[2(w_1^{(t)^2} - w_2^{(t)})w_1^{(t)} + w_1^{(t)} - 1 \right]$$

$$w_2^{(t+1)} \leftarrow w_2^{(t)} - \eta \left[-(w_1^{(t)^2} - w_2^{(t)}) \right]$$

$$t \leftarrow t + 1$$

An example

Example: $F(\mathbf{w}) = 0.5(w_1^2 - w_2)^2 + 0.5(w_1 - 1)^2$. Gradient is

$$\frac{\partial F}{\partial w_1} = 2(w_1^2 - w_2)w_1 + w_1 - 1 \quad \frac{\partial F}{\partial w_2} = -(w_1^2 - w_2)$$

GD:

- Initialize $w_1^{(0)}$ and $w_2^{(0)}$ (to be 0 or *randomly*), $t = 0$
- do

$$w_1^{(t+1)} \leftarrow w_1^{(t)} - \eta \left[2(w_1^{(t)^2} - w_2^{(t)})w_1^{(t)} + w_1^{(t)} - 1 \right]$$

$$w_2^{(t+1)} \leftarrow w_2^{(t)} - \eta \left[-(w_1^{(t)^2} - w_2^{(t)}) \right]$$

$$t \leftarrow t + 1$$

- until $F(\mathbf{w}^{(t)})$ **does not change much**

Why GD?

Intuition: by first-order **Taylor approximation**

$$F(\boldsymbol{w}) \approx F(\boldsymbol{w}^{(t)}) + \nabla F(\boldsymbol{w}^{(t)})^T (\boldsymbol{w} - \boldsymbol{w}^{(t)})$$

Why GD?

Intuition: by first-order **Taylor approximation**

$$F(\mathbf{w}) \approx F(\mathbf{w}^{(t)}) + \nabla F(\mathbf{w}^{(t)})^T (\mathbf{w} - \mathbf{w}^{(t)})$$

GD ensures

$$F(\mathbf{w}^{(t+1)}) \approx F(\mathbf{w}^{(t)}) - \eta \|\nabla F(\mathbf{w}^{(t)})\|_2^2 \leq F(\mathbf{w}^{(t)})$$

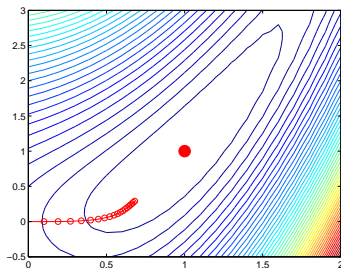
Why GD?

Intuition: by first-order **Taylor approximation**

$$F(\mathbf{w}) \approx F(\mathbf{w}^{(t)}) + \nabla F(\mathbf{w}^{(t)})^T (\mathbf{w} - \mathbf{w}^{(t)})$$

GD ensures

$$F(\mathbf{w}^{(t+1)}) \approx F(\mathbf{w}^{(t)}) - \eta \|\nabla F(\mathbf{w}^{(t)})\|_2^2 \leq F(\mathbf{w}^{(t)})$$



reasonable η decreases function value

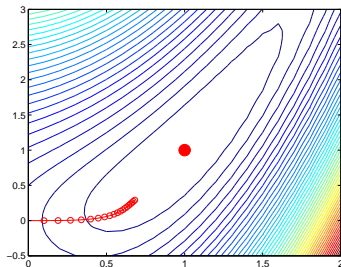
Why GD?

Intuition: by first-order **Taylor approximation**

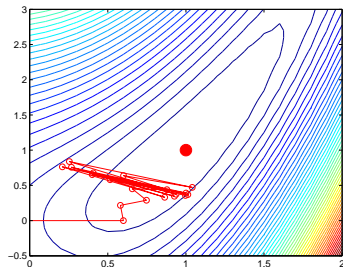
$$F(\mathbf{w}) \approx F(\mathbf{w}^{(t)}) + \nabla F(\mathbf{w}^{(t)})^T (\mathbf{w} - \mathbf{w}^{(t)})$$

GD ensures

$$F(\mathbf{w}^{(t+1)}) \approx F(\mathbf{w}^{(t)}) - \eta \|\nabla F(\mathbf{w}^{(t)})\|_2^2 \leq F(\mathbf{w}^{(t)})$$



reasonable η decreases function value



but large η is unstable

Stochastic Gradient Descent (SGD)

GD: keep moving in the negative gradient direction

SGD: keep moving in some *noisy* negative gradient direction

Stochastic Gradient Descent (SGD)

GD: keep moving in the negative gradient direction

SGD: keep moving in some *noisy* negative gradient direction

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \eta \tilde{\nabla} F(\mathbf{w}^{(t)})$$

where $\tilde{\nabla} F(\mathbf{w}^{(t)})$ is a random variable (called **stochastic gradient**) s.t.

$$\mathbb{E} \left[\tilde{\nabla} F(\mathbf{w}^{(t)}) \right] = \nabla F(\mathbf{w}^{(t)}) \quad (\text{unbiasedness})$$

Stochastic Gradient Descent (SGD)

GD: keep moving in the negative gradient direction

SGD: keep moving in some *noisy* negative gradient direction

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \eta \tilde{\nabla} F(\mathbf{w}^{(t)})$$

where $\tilde{\nabla} F(\mathbf{w}^{(t)})$ is a random variable (called **stochastic gradient**) s.t.

$$\mathbb{E} \left[\tilde{\nabla} F(\mathbf{w}^{(t)}) \right] = \nabla F(\mathbf{w}^{(t)}) \quad (\text{unbiasedness})$$

Key point: it could be *much faster to obtain a stochastic gradient!*

Convergence Guarantees

Many for both GD and SGD on convex objectives.

Convergence Guarantees

Many for both GD and SGD on convex objectives.

They tell you at most how many iterations you need to achieve

$$F(\mathbf{w}^{(t)}) - F(\mathbf{w}^*) \leq \epsilon$$

Convergence Guarantees

Many for both GD and SGD on convex objectives.

They tell you at most how many iterations you need to achieve

$$F(\mathbf{w}^{(t)}) - F(\mathbf{w}^*) \leq \epsilon$$

Even for *nonconvex objectives*, many recent works show effectiveness of GD/SGD.

Applying GD to perceptron loss

Objective

$$F(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N \max\{0, -y_n \mathbf{w}^T \mathbf{x}_n\}$$

Applying GD to perceptron loss

Objective

$$F(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N \max\{0, -y_n \mathbf{w}^T \mathbf{x}_n\}$$

Gradient (or really *sub-gradient*) is

$$\nabla F(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N -\mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

(only misclassified examples contribute to the gradient)

Applying GD to perceptron loss

Objective

$$F(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N \max\{0, -y_n \mathbf{w}^T \mathbf{x}_n\}$$

Gradient (or really *sub-gradient*) is

$$\nabla F(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N -\mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

(only misclassified examples contribute to the gradient)

GD update

$$\mathbf{w} \leftarrow \mathbf{w} + \frac{\eta}{N} \sum_{n=1}^N \mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

Applying GD to perceptron loss

Objective

$$F(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N \max\{0, -y_n \mathbf{w}^T \mathbf{x}_n\}$$

Gradient (or really *sub-gradient*) is

$$\nabla F(\mathbf{w}) = \frac{1}{N} \sum_{n=1}^N -\mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

(only misclassified examples contribute to the gradient)

GD update

$$\mathbf{w} \leftarrow \mathbf{w} + \frac{\eta}{N} \sum_{n=1}^N \mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

Slow: each update makes one pass of the entire training set!

Applying SGD to perceptron loss

How to construct a stochastic gradient?

Applying SGD to perceptron loss

How to construct a stochastic gradient?

One common trick: pick one example $n \in [N]$ uniformly at random, let

$$\tilde{\nabla} F(\mathbf{w}^{(t)}) = -\mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

clearly unbiased (convince yourself).

Applying SGD to perceptron loss

How to construct a stochastic gradient?

One common trick: pick one example $n \in [N]$ uniformly at random, let

$$\tilde{\nabla} F(\mathbf{w}^{(t)}) = -\mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

clearly unbiased (convince yourself).

SGD update:

$$\mathbf{w} \leftarrow \mathbf{w} + \eta \mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

Applying SGD to perceptron loss

How to construct a stochastic gradient?

One common trick: pick one example $n \in [N]$ uniformly at random, let

$$\tilde{\nabla} F(\mathbf{w}^{(t)}) = -\mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

clearly unbiased (convince yourself).

SGD update:

$$\mathbf{w} \leftarrow \mathbf{w} + \eta \mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

Fast: each update touches only one data point!

Applying SGD to perceptron loss

How to construct a stochastic gradient?

One common trick: pick one example $n \in [N]$ uniformly at random, let

$$\tilde{\nabla} F(\mathbf{w}^{(t)}) = -\mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

clearly unbiased (convince yourself).

SGD update:

$$\mathbf{w} \leftarrow \mathbf{w} + \eta \mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

Fast: each update touches only one data point!

Conveniently, objective of most ML tasks is a *finite sum* (over each training point) and the above trick applies!

Applying SGD to perceptron loss

How to construct a stochastic gradient?

One common trick: pick one example $n \in [N]$ uniformly at random, let

$$\tilde{\nabla} F(\mathbf{w}^{(t)}) = -\mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

clearly unbiased (convince yourself).

SGD update:

$$\mathbf{w} \leftarrow \mathbf{w} + \eta \mathbb{I}[y_n \mathbf{w}^T \mathbf{x}_n \leq 0] y_n \mathbf{x}_n$$

Fast: each update touches only one data point!

Conveniently, objective of most ML tasks is a *finite sum* (over each training point) and the above trick applies!

Exercise: try SGD to minimize RSS for linear regression.

The Perceptron Algorithm

Perceptron algorithm is SGD with $\eta = 1$ applied to perceptron loss:

The Perceptron Algorithm

Perceptron algorithm is SGD with $\eta = 1$ applied to perceptron loss:

Repeat:

- Pick a data point \mathbf{x}_n uniformly at random
- If $\text{sgn}(\mathbf{w}^T \mathbf{x}_n) \neq y_n$

$$\mathbf{w} \leftarrow \mathbf{w} + y_n \mathbf{x}_n$$

The Perceptron Algorithm

Perceptron algorithm is SGD with $\eta = 1$ applied to perceptron loss:

Repeat:

- Pick a data point \mathbf{x}_n uniformly at random
- If $\text{sgn}(\mathbf{w}^T \mathbf{x}_n) \neq y_n$

$$\mathbf{w} \leftarrow \mathbf{w} + y_n \mathbf{x}_n$$

Note:

- \mathbf{w} is always a *linear combination* of the training examples

The Perceptron Algorithm

Perceptron algorithm is SGD with $\eta = 1$ applied to perceptron loss:

Repeat:

- Pick a data point \mathbf{x}_n uniformly at random
- If $\text{sgn}(\mathbf{w}^T \mathbf{x}_n) \neq y_n$

$$\mathbf{w} \leftarrow \mathbf{w} + y_n \mathbf{x}_n$$

Note:

- \mathbf{w} is always a *linear combination* of the training examples
- why $\eta = 1$? Does not really matter in terms of training error

Why does it make sense?

If the current weight \mathbf{w} makes a mistake

$$y_n \mathbf{w}^T \mathbf{x}_n < 0$$

Why does it make sense?

If the current weight \mathbf{w} makes a mistake

$$y_n \mathbf{w}^T \mathbf{x}_n < 0$$

then after the update $\mathbf{w}' = \mathbf{w} + y_n \mathbf{x}_n$ we have

$$y_n \mathbf{w}'^T \mathbf{x}_n = y_n \mathbf{w}^T \mathbf{x}_n + y_n^2 \mathbf{x}_n^T \mathbf{x}_n \geq y_n \mathbf{w}^T \mathbf{x}_n$$

Why does it make sense?

If the current weight \mathbf{w} makes a mistake

$$y_n \mathbf{w}^T \mathbf{x}_n < 0$$

then after the update $\mathbf{w}' = \mathbf{w} + y_n \mathbf{x}_n$ we have

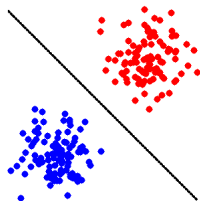
$$y_n \mathbf{w}'^T \mathbf{x}_n = y_n \mathbf{w}^T \mathbf{x}_n + y_n^2 \mathbf{x}_n^T \mathbf{x}_n \geq y_n \mathbf{w}^T \mathbf{x}_n$$

Thus it is more likely to get it right after the update.

Any theory?

(HW 1) If training set is linearly separable

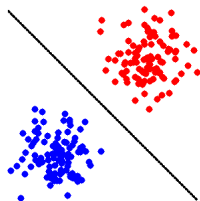
- Perceptron *converges in a finite number of steps*
- training error is 0



Any theory?

(HW 1) If training set is linearly separable

- Perceptron *converges in a finite number of steps*
- training error is 0



There are also guarantees when the data are not linearly separable.

Outline

- 1 Review of Last Lecture
- 2 Linear Classifier and Surrogate Losses
- 3 Perceptron
- 4 **Logistic regression**
 - A Probabilistic View
 - Optimization

A simple view

In one sentence: find the minimizer of

$$\begin{aligned} F(\mathbf{w}) &= \frac{1}{N} \sum_{n=1}^N \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) \\ &= \frac{1}{N} \sum_{n=1}^N \ln(1 + e^{-y_n \mathbf{w}^T \mathbf{x}_n}) \end{aligned}$$

A simple view

In one sentence: find the minimizer of

$$\begin{aligned} F(\mathbf{w}) &= \frac{1}{N} \sum_{n=1}^N \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) \\ &= \frac{1}{N} \sum_{n=1}^N \ln(1 + e^{-y_n \mathbf{w}^T \mathbf{x}_n}) \end{aligned}$$

But why logistic loss? and why “regression”?

Predicting probability

Instead of predicting a discrete label, can we *predict the probability of each label?* i.e. regress the probabilities

Predicting probability

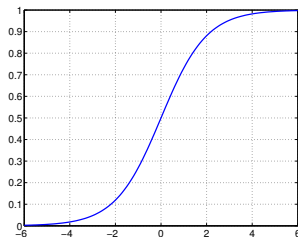
Instead of predicting a discrete label, can we *predict the probability of each label?* i.e. regress the probabilities

One way: **sigmoid function + linear model**

$$\mathbb{P}(y = +1 \mid \mathbf{x}; \mathbf{w}) = \sigma(\mathbf{w}^T \mathbf{x})$$

where σ is the sigmoid function:

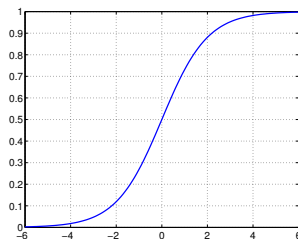
$$\sigma(z) = \frac{1}{1 + e^{-z}}$$



Properties

Properties of sigmoid $\sigma(z) = \frac{1}{1+e^{-z}}$

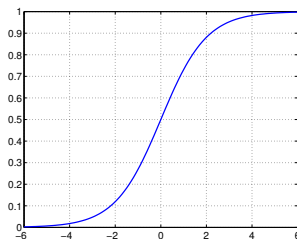
- between 0 and 1 (good as probability)



Properties

Properties of sigmoid $\sigma(z) = \frac{1}{1+e^{-z}}$

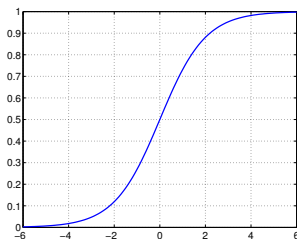
- between 0 and 1 (good as probability)
- $\sigma(\mathbf{w}^T \mathbf{x}) \geq 0.5 \Leftrightarrow \mathbf{w}^T \mathbf{x} \geq 0$, consistent with predicting the label with $\text{sgn}(\mathbf{w}^T \mathbf{x})$



Properties

Properties of sigmoid $\sigma(z) = \frac{1}{1+e^{-z}}$

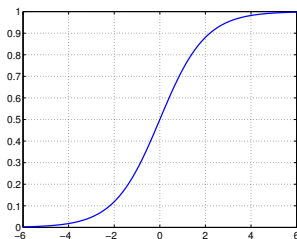
- between 0 and 1 (good as probability)
- $\sigma(\mathbf{w}^T \mathbf{x}) \geq 0.5 \Leftrightarrow \mathbf{w}^T \mathbf{x} \geq 0$, consistent with predicting the label with $\text{sgn}(\mathbf{w}^T \mathbf{x})$
- larger $\mathbf{w}^T \mathbf{x} \Rightarrow$ larger $\sigma(\mathbf{w}^T \mathbf{x}) \Rightarrow$ higher *confidence* in label 1



Properties

Properties of sigmoid $\sigma(z) = \frac{1}{1+e^{-z}}$

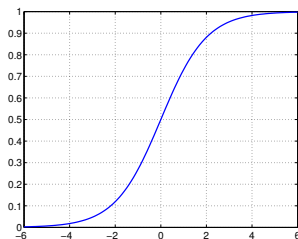
- between 0 and 1 (good as probability)
- $\sigma(\mathbf{w}^T \mathbf{x}) \geq 0.5 \Leftrightarrow \mathbf{w}^T \mathbf{x} \geq 0$, consistent with predicting the label with $\text{sgn}(\mathbf{w}^T \mathbf{x})$
- larger $\mathbf{w}^T \mathbf{x} \Rightarrow$ larger $\sigma(\mathbf{w}^T \mathbf{x}) \Rightarrow$ higher *confidence* in label 1
- $\sigma(z) + \sigma(-z) = 1$ for all z



Properties

Properties of sigmoid $\sigma(z) = \frac{1}{1+e^{-z}}$

- between 0 and 1 (good as probability)
- $\sigma(\mathbf{w}^T \mathbf{x}) \geq 0.5 \Leftrightarrow \mathbf{w}^T \mathbf{x} \geq 0$, consistent with predicting the label with $\text{sgn}(\mathbf{w}^T \mathbf{x})$
- larger $\mathbf{w}^T \mathbf{x} \Rightarrow$ larger $\sigma(\mathbf{w}^T \mathbf{x}) \Rightarrow$ higher *confidence* in label 1
- $\sigma(z) + \sigma(-z) = 1$ for all z



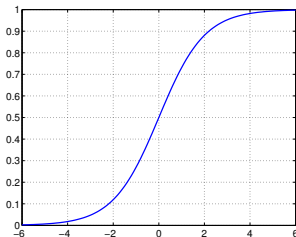
The probability of label -1 is naturally

$$1 - \mathbb{P}(y = +1 \mid \mathbf{x}; \mathbf{w}) = 1 - \sigma(\mathbf{w}^T \mathbf{x}) = \sigma(-\mathbf{w}^T \mathbf{x})$$

Properties

Properties of sigmoid $\sigma(z) = \frac{1}{1+e^{-z}}$

- between 0 and 1 (good as probability)
- $\sigma(\mathbf{w}^T \mathbf{x}) \geq 0.5 \Leftrightarrow \mathbf{w}^T \mathbf{x} \geq 0$, consistent with predicting the label with $\text{sgn}(\mathbf{w}^T \mathbf{x})$
- larger $\mathbf{w}^T \mathbf{x} \Rightarrow$ larger $\sigma(\mathbf{w}^T \mathbf{x}) \Rightarrow$ higher *confidence* in label 1
- $\sigma(z) + \sigma(-z) = 1$ for all z



The probability of label -1 is naturally

$$1 - \mathbb{P}(y = +1 \mid \mathbf{x}; \mathbf{w}) = 1 - \sigma(\mathbf{w}^T \mathbf{x}) = \sigma(-\mathbf{w}^T \mathbf{x})$$

and thus

$$\mathbb{P}(y \mid \mathbf{x}; \mathbf{w}) = \sigma(y\mathbf{w}^T \mathbf{x}) = \frac{1}{1 + e^{-y\mathbf{w}^T \mathbf{x}}}$$

How to regress with discrete labels?

What we observe are labels, not probabilities.

How to regress with discrete labels?

What we observe are labels, not probabilities.

Take a **probabilistic view**

- assume data is generated in this way by some w
- perform Maximum Likelihood Estimation (MLE)

How to regress with discrete labels?

What we observe are labels, not probabilities.

Take a **probabilistic view**

- assume data is generated in this way by some w
- perform Maximum Likelihood Estimation (MLE)

Specifically, what is the probability of seeing label y_1, \dots, y_n given x_1, \dots, x_n , as a function of some w ?

$$P(w) = \prod_{n=1}^N \mathbb{P}(y_n \mid \mathbf{x}_n; w)$$

MLE: find w^* that **maximizes the probability** $P(w)$

The MLE solution

$$\boldsymbol{w}^* = \operatorname{argmax}_{\boldsymbol{w}} P(\boldsymbol{w}) = \operatorname{argmax}_{\boldsymbol{w}} \prod_{n=1}^N \mathbb{P}(y_n \mid \boldsymbol{x}_n; \boldsymbol{w})$$

The MLE solution

$$\begin{aligned}\mathbf{w}^* &= \operatorname{argmax}_{\mathbf{w}} P(\mathbf{w}) = \operatorname{argmax}_{\mathbf{w}} \prod_{n=1}^N \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) \\ &= \operatorname{argmax}_{\mathbf{w}} \sum_{n=1}^N \ln \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w})\end{aligned}$$

The MLE solution

$$\begin{aligned}\mathbf{w}^* &= \operatorname{argmax}_{\mathbf{w}} P(\mathbf{w}) = \operatorname{argmax}_{\mathbf{w}} \prod_{n=1}^N \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) \\ &= \operatorname{argmax}_{\mathbf{w}} \sum_{n=1}^N \ln \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) = \operatorname{argmin}_{\mathbf{w}} \sum_{n=1}^N -\ln \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w})\end{aligned}$$

The MLE solution

$$\begin{aligned}\mathbf{w}^* &= \operatorname{argmax}_{\mathbf{w}} P(\mathbf{w}) = \operatorname{argmax}_{\mathbf{w}} \prod_{n=1}^N \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) \\ &= \operatorname{argmax}_{\mathbf{w}} \sum_{n=1}^N \ln \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) = \operatorname{argmin}_{\mathbf{w}} \sum_{n=1}^N -\ln \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) \\ &= \operatorname{argmin}_{\mathbf{w}} \sum_{n=1}^N \ln(1 + e^{-y_n \mathbf{w}^\top \mathbf{x}_n})\end{aligned}$$

The MLE solution

$$\begin{aligned}\mathbf{w}^* &= \operatorname{argmax}_{\mathbf{w}} P(\mathbf{w}) = \operatorname{argmax}_{\mathbf{w}} \prod_{n=1}^N \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) \\ &= \operatorname{argmax}_{\mathbf{w}} \sum_{n=1}^N \ln \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) = \operatorname{argmin}_{\mathbf{w}} \sum_{n=1}^N -\ln \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) \\ &= \operatorname{argmin}_{\mathbf{w}} \sum_{n=1}^N \ln(1 + e^{-y_n \mathbf{w}^\top \mathbf{x}_n}) = \operatorname{argmin}_{\mathbf{w}} \sum_{n=1}^N \ell_{\text{logistic}}(y_n \mathbf{w}^\top \mathbf{x}_n)\end{aligned}$$

The MLE solution

$$\begin{aligned}\mathbf{w}^* &= \operatorname{argmax}_{\mathbf{w}} P(\mathbf{w}) = \operatorname{argmax}_{\mathbf{w}} \prod_{n=1}^N \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) \\ &= \operatorname{argmax}_{\mathbf{w}} \sum_{n=1}^N \ln \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) = \operatorname{argmin}_{\mathbf{w}} \sum_{n=1}^N -\ln \mathbb{P}(y_n \mid \mathbf{x}_n; \mathbf{w}) \\ &= \operatorname{argmin}_{\mathbf{w}} \sum_{n=1}^N \ln(1 + e^{-y_n \mathbf{w}^T \mathbf{x}_n}) = \operatorname{argmin}_{\mathbf{w}} \sum_{n=1}^N \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) \\ &= \operatorname{argmin}_{\mathbf{w}} F(\mathbf{w})\end{aligned}$$

i.e. *minimizing logistic loss is exactly doing MLE for the sigmoid model!*

Let's apply SGD again

$$\boldsymbol{w} \leftarrow \boldsymbol{w} - \eta \tilde{\nabla} F(\boldsymbol{w})$$

Let's apply SGD again

$$\begin{aligned}\mathbf{w} &\leftarrow \mathbf{w} - \eta \tilde{\nabla} F(\mathbf{w}) \\ &= \mathbf{w} - \eta \nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) \quad (n \in [N] \text{ is drawn u.a.r.})\end{aligned}$$

Let's apply SGD again

$$\begin{aligned}\mathbf{w} &\leftarrow \mathbf{w} - \eta \tilde{\nabla} F(\mathbf{w}) \\ &= \mathbf{w} - \eta \nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) \quad (n \in [N] \text{ is drawn u.a.r.}) \\ &= \mathbf{w} - \eta \left(\frac{\partial \ell_{\text{logistic}}(z)}{\partial z} \Big|_{z=y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n \mathbf{x}_n\end{aligned}$$

Let's apply SGD again

$$\begin{aligned}\mathbf{w} &\leftarrow \mathbf{w} - \eta \tilde{\nabla} F(\mathbf{w}) \\ &= \mathbf{w} - \eta \nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) \quad (n \in [N] \text{ is drawn u.a.r.}) \\ &= \mathbf{w} - \eta \left(\frac{\partial \ell_{\text{logistic}}(z)}{\partial z} \Big|_{z=y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n \mathbf{x}_n \\ &= \mathbf{w} - \eta \left(\frac{-e^{-z}}{1 + e^{-z}} \Big|_{z=y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n \mathbf{x}_n\end{aligned}$$

Let's apply SGD again

$$\begin{aligned}\mathbf{w} &\leftarrow \mathbf{w} - \eta \tilde{\nabla} F(\mathbf{w}) \\ &= \mathbf{w} - \eta \nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) \quad (n \in [N] \text{ is drawn u.a.r.}) \\ &= \mathbf{w} - \eta \left(\frac{\partial \ell_{\text{logistic}}(z)}{\partial z} \Big|_{z=y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n \mathbf{x}_n \\ &= \mathbf{w} - \eta \left(\frac{-e^{-z}}{1 + e^{-z}} \Big|_{z=y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n \mathbf{x}_n \\ &= \mathbf{w} + \eta \sigma(-y_n \mathbf{w}^T \mathbf{x}_n) y_n \mathbf{x}_n\end{aligned}$$

Let's apply SGD again

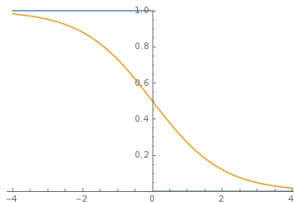
$$\begin{aligned}\mathbf{w} &\leftarrow \mathbf{w} - \eta \tilde{\nabla} F(\mathbf{w}) \\&= \mathbf{w} - \eta \nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) \quad (n \in [N] \text{ is drawn u.a.r.}) \\&= \mathbf{w} - \eta \left(\frac{\partial \ell_{\text{logistic}}(z)}{\partial z} \Big|_{z=y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n \mathbf{x}_n \\&= \mathbf{w} - \eta \left(\frac{-e^{-z}}{1 + e^{-z}} \Big|_{z=y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n \mathbf{x}_n \\&= \mathbf{w} + \eta \sigma(-y_n \mathbf{w}^T \mathbf{x}_n) y_n \mathbf{x}_n \\&= \mathbf{w} + \eta \mathbb{P}(-y_n \mid \mathbf{x}_n; \mathbf{w}) y_n \mathbf{x}_n\end{aligned}$$

Let's apply SGD again

$$\begin{aligned}
 \mathbf{w} &\leftarrow \mathbf{w} - \eta \tilde{\nabla} F(\mathbf{w}) \\
 &= \mathbf{w} - \eta \nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) \quad (n \in [N] \text{ is drawn u.a.r.}) \\
 &= \mathbf{w} - \eta \left(\frac{\partial \ell_{\text{logistic}}(z)}{\partial z} \Big|_{z=y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n \mathbf{x}_n \\
 &= \mathbf{w} - \eta \left(\frac{-e^{-z}}{1 + e^{-z}} \Big|_{z=y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n \mathbf{x}_n \\
 &= \mathbf{w} + \eta \sigma(-y_n \mathbf{w}^T \mathbf{x}_n) y_n \mathbf{x}_n \\
 &= \mathbf{w} + \eta \mathbb{P}(-y_n \mid \mathbf{x}_n; \mathbf{w}) y_n \mathbf{x}_n
 \end{aligned}$$

This is a *soft version of Perceptron!*

$\mathbb{P}(-y_n \mid \mathbf{x}_n; \mathbf{w})$ versus $\mathbb{I}[y_n \neq \text{sgn}(\mathbf{w}^T \mathbf{x}_n)]$



A second-order method: Newton method

Recall the intuition of GD: we look at first-order **Taylor approximation**

$$F(\mathbf{w}) \approx F(\mathbf{w}^{(t)}) + \nabla F(\mathbf{w}^{(t)})^T (\mathbf{w} - \mathbf{w}^{(t)})$$

A second-order method: Newton method

Recall the intuition of GD: we look at first-order **Taylor approximation**

$$F(\mathbf{w}) \approx F(\mathbf{w}^{(t)}) + \nabla F(\mathbf{w}^{(t)})^T (\mathbf{w} - \mathbf{w}^{(t)})$$

What if we look at *second-order* Taylor approximation?

$$F(\mathbf{w}) \approx F(\mathbf{w}^{(t)}) + \nabla F(\mathbf{w}^{(t)})^T (\mathbf{w} - \mathbf{w}^{(t)}) + \frac{1}{2} (\mathbf{w} - \mathbf{w}^{(t)})^T \mathbf{H}_t (\mathbf{w} - \mathbf{w}^{(t)})$$

A second-order method: Newton method

Recall the intuition of GD: we look at first-order **Taylor approximation**

$$F(\mathbf{w}) \approx F(\mathbf{w}^{(t)}) + \nabla F(\mathbf{w}^{(t)})^T (\mathbf{w} - \mathbf{w}^{(t)})$$

What if we look at *second-order* Taylor approximation?

$$F(\mathbf{w}) \approx F(\mathbf{w}^{(t)}) + \nabla F(\mathbf{w}^{(t)})^T (\mathbf{w} - \mathbf{w}^{(t)}) + \frac{1}{2} (\mathbf{w} - \mathbf{w}^{(t)})^T \mathbf{H}_t (\mathbf{w} - \mathbf{w}^{(t)})$$

where $\mathbf{H}_t = \nabla^2 F(\mathbf{w}^{(t)}) \in \mathbb{R}^{D \times D}$ is the *Hessian* of F at $\mathbf{w}^{(t)}$, i.e.,

$$H_{t,ij} = \left. \frac{\partial^2 F(\mathbf{w})}{\partial w_i \partial w_j} \right|_{\mathbf{w}=\mathbf{w}^{(t)}}$$

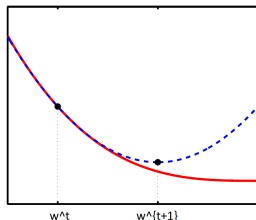
(think “second derivative” when $D = 1$)

Deriving Newton method

If we minimize the second-order approximation (via “complete the square”)

$$F(\mathbf{w})$$

$$\begin{aligned} &\approx F(\mathbf{w}^{(t)}) + \nabla F(\mathbf{w}^{(t)})^T (\mathbf{w} - \mathbf{w}^{(t)}) + \frac{1}{2} (\mathbf{w} - \mathbf{w}^{(t)})^T \mathbf{H}_t (\mathbf{w} - \mathbf{w}^{(t)}) \\ &= \frac{1}{2} \left(\mathbf{w} - \mathbf{w}^{(t)} + \mathbf{H}_t^{-1} \nabla F(\mathbf{w}^{(t)}) \right)^T \mathbf{H}_t \left(\mathbf{w} - \mathbf{w}^{(t)} + \mathbf{H}_t^{-1} \nabla F(\mathbf{w}^{(t)}) \right) + \text{cnt.} \end{aligned}$$



Deriving Newton method

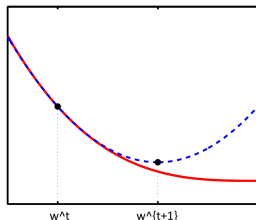
If we minimize the second-order approximation (via “complete the square”)

$$F(\mathbf{w})$$

$$\begin{aligned} &\approx F(\mathbf{w}^{(t)}) + \nabla F(\mathbf{w}^{(t)})^T (\mathbf{w} - \mathbf{w}^{(t)}) + \frac{1}{2} (\mathbf{w} - \mathbf{w}^{(t)})^T \mathbf{H}_t (\mathbf{w} - \mathbf{w}^{(t)}) \\ &= \frac{1}{2} \left(\mathbf{w} - \mathbf{w}^{(t)} + \mathbf{H}_t^{-1} \nabla F(\mathbf{w}^{(t)}) \right)^T \mathbf{H}_t \left(\mathbf{w} - \mathbf{w}^{(t)} + \mathbf{H}_t^{-1} \nabla F(\mathbf{w}^{(t)}) \right) + \text{cnt.} \end{aligned}$$

for convex F (so H_t is *positive semidefinite*)
we obtain **Newton method**:

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \mathbf{H}_t^{-1} \nabla F(\mathbf{w}^{(t)})$$



Comparing GD and Newton

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \eta \nabla F(\mathbf{w}^{(t)}) \quad (\text{GD})$$

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \mathbf{H}_t^{-1} \nabla F(\mathbf{w}^{(t)}) \quad (\text{Newton})$$

Both are iterative optimization procedures,

Comparing GD and Newton

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \eta \nabla F(\mathbf{w}^{(t)}) \quad (\text{GD})$$

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \mathbf{H}_t^{-1} \nabla F(\mathbf{w}^{(t)}) \quad (\text{Newton})$$

Both are iterative optimization procedures, but Newton method

- has no learning rate η (*so no tuning needed!*)

Comparing GD and Newton

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \eta \nabla F(\mathbf{w}^{(t)}) \quad (\text{GD})$$

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \mathbf{H}_t^{-1} \nabla F(\mathbf{w}^{(t)}) \quad (\text{Newton})$$

Both are iterative optimization procedures, but Newton method

- has no learning rate η (*so no tuning needed!*)
- converges *super fast* in terms of #iterations needed

Comparing GD and Newton

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \eta \nabla F(\mathbf{w}^{(t)}) \quad (\text{GD})$$

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \mathbf{H}_t^{-1} \nabla F(\mathbf{w}^{(t)}) \quad (\text{Newton})$$

Both are iterative optimization procedures, but Newton method

- has no learning rate η (*so no tuning needed!*)
- converges *super fast* in terms of #iterations needed
 - e.g. how many iterations needed when applied to a quadratic?

Comparing GD and Newton

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \eta \nabla F(\mathbf{w}^{(t)}) \quad (\text{GD})$$

$$\mathbf{w}^{(t+1)} \leftarrow \mathbf{w}^{(t)} - \mathbf{H}_t^{-1} \nabla F(\mathbf{w}^{(t)}) \quad (\text{Newton})$$

Both are iterative optimization procedures, but Newton method

- has no learning rate η (*so no tuning needed!*)
- converges *super fast* in terms of #iterations needed
 - e.g. how many iterations needed when applied to a quadratic?
- requires **second-order** information and is *slow* each iteration (there are many ways to improve it though)

Applying Newton to logistic loss

$$\nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) = -\sigma(-y_n \mathbf{w}^T \mathbf{x}_n) y_n \mathbf{x}_n$$

Applying Newton to logistic loss

$$\nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) = -\sigma(-y_n \mathbf{w}^T \mathbf{x}_n) y_n \mathbf{x}_n$$

$$\nabla_{\mathbf{w}}^2 \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) = \left(\frac{\partial \sigma(z)}{\partial z} \Big|_{z=-y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n^2 \mathbf{x}_n \mathbf{x}_n^T$$

Applying Newton to logistic loss

$$\nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) = -\sigma(-y_n \mathbf{w}^T \mathbf{x}_n) y_n \mathbf{x}_n$$

$$\begin{aligned} \nabla_{\mathbf{w}}^2 \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) &= \left(\frac{\partial \sigma(z)}{\partial z} \Big|_{z=-y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n^2 \mathbf{x}_n \mathbf{x}_n^T \\ &= \left(\frac{e^{-z}}{(1 + e^{-z})^2} \Big|_{z=-y_n \mathbf{w}^T \mathbf{x}_n} \right) \mathbf{x}_n \mathbf{x}_n^T \end{aligned}$$

Applying Newton to logistic loss

$$\nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) = -\sigma(-y_n \mathbf{w}^T \mathbf{x}_n) y_n \mathbf{x}_n$$

$$\begin{aligned} \nabla_{\mathbf{w}}^2 \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) &= \left(\frac{\partial \sigma(z)}{\partial z} \Big|_{z=-y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n^2 \mathbf{x}_n \mathbf{x}_n^T \\ &= \left(\frac{e^{-z}}{(1 + e^{-z})^2} \Big|_{z=-y_n \mathbf{w}^T \mathbf{x}_n} \right) \mathbf{x}_n \mathbf{x}_n^T \\ &= \sigma(y_n \mathbf{w}^T \mathbf{x}_n) (1 - \sigma(y_n \mathbf{w}^T \mathbf{x}_n)) \mathbf{x}_n \mathbf{x}_n^T \end{aligned}$$

Applying Newton to logistic loss

$$\nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) = -\sigma(-y_n \mathbf{w}^T \mathbf{x}_n) y_n \mathbf{x}_n$$

$$\begin{aligned} \nabla_{\mathbf{w}}^2 \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) &= \left(\frac{\partial \sigma(z)}{\partial z} \Big|_{z=-y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n^2 \mathbf{x}_n \mathbf{x}_n^T \\ &= \left(\frac{e^{-z}}{(1 + e^{-z})^2} \Big|_{z=-y_n \mathbf{w}^T \mathbf{x}_n} \right) \mathbf{x}_n \mathbf{x}_n^T \\ &= \sigma(y_n \mathbf{w}^T \mathbf{x}_n) (1 - \sigma(y_n \mathbf{w}^T \mathbf{x}_n)) \mathbf{x}_n \mathbf{x}_n^T \end{aligned}$$

Exercises:

- why is the Hessian of logistic loss positive semidefinite?

Applying Newton to logistic loss

$$\nabla_{\mathbf{w}} \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) = -\sigma(-y_n \mathbf{w}^T \mathbf{x}_n) y_n \mathbf{x}_n$$

$$\begin{aligned} \nabla_{\mathbf{w}}^2 \ell_{\text{logistic}}(y_n \mathbf{w}^T \mathbf{x}_n) &= \left(\frac{\partial \sigma(z)}{\partial z} \Big|_{z=-y_n \mathbf{w}^T \mathbf{x}_n} \right) y_n^2 \mathbf{x}_n \mathbf{x}_n^T \\ &= \left(\frac{e^{-z}}{(1 + e^{-z})^2} \Big|_{z=-y_n \mathbf{w}^T \mathbf{x}_n} \right) \mathbf{x}_n \mathbf{x}_n^T \\ &= \sigma(y_n \mathbf{w}^T \mathbf{x}_n) (1 - \sigma(y_n \mathbf{w}^T \mathbf{x}_n)) \mathbf{x}_n \mathbf{x}_n^T \end{aligned}$$

Exercises:

- why is the Hessian of logistic loss positive semidefinite?
- can we apply Newton method to perceptron/hinge loss?

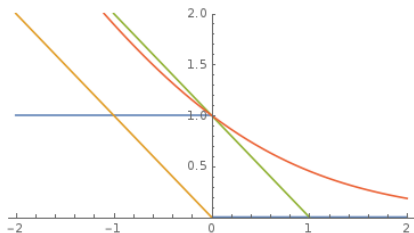
Summary

Linear models for classification:

Step 1. Model is the set of **separating hyperplanes**

$$\mathcal{F} = \{f(\mathbf{x}) = \text{sgn}(\mathbf{w}^T \mathbf{x}) \mid \mathbf{w} \in \mathbb{R}^D\}$$

Step 2. Pick the **surrogate loss**



- **perceptron loss** $\ell_{\text{perceptron}}(z) = \max\{0, -z\}$ (used in Perceptron)
- **hinge loss** $\ell_{\text{hinge}}(z) = \max\{0, 1 - z\}$ (used in SVM and many others)
- **logistic loss** $\ell_{\text{logistic}}(z) = \log(1 + \exp(-z))$ (used in logistic regression)

Step 3. Find empirical risk minimizer (ERM):

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^D} \frac{1}{N} \sum_{n=1}^N \ell(y_n \mathbf{w}^T \mathbf{x}_n)$$

using **GD/SGD/Newton**.