

DSC 140A

Probabilistic Modeling & Machine Learning

Lecture 6 | Part 1


Linear Classification

Classification

- ▶ We've been talking about **regression**.
 - ▶ Label is a continuous value.
- ▶ What about **classification**?
 - ▶ Label is a discrete value.

Example: Penguins

- Given a new penguin's measurements, predict its species.



Looking Back

- ▶ We know one classification algorithm already.
 - ▶ k -Nearest Neighbors.
- ▶ But k -NN does not “**learn**”, it “**memorizes**”.
- ▶ Can we use linear predictors for classification?

$$H(\vec{X}) = w_0 + w_1(\text{flipper length}) + w_2(\text{body mass})$$

- ▶ Train by minimizing risk?

Linear Classifiers

- ▶ **Problem:** output of $H(\vec{x})$ is a real number; we want the output to be a **species**.

$$H(\vec{x}) = w_0 + w_1(\text{flipper length}) + w_2(\text{body mass})$$

- ▶ **Idea:** turn species into a number.

Label Encodings

- ▶ There are two natural ways to **encode** a label y as a number in **binary classification**.
- ▶ $y \in \{0, 1\}$:
 - ▶ $y = 0$ for one class, $y = 1$ for the other.
 - ▶ **Example:** 0 for Adelie, 1 for Gentoo.
- ▶ $y \in \{-1, 1\}$:
 - ▶ $y = -1$ for one class, $y = 1$ for the other.
 - ▶ **Example:** -1 for Adelie, 1 for Gentoo.
- ▶ Unless otherwise specified, we'll use $y \in \{-1, 1\}$.

Linear Classifiers

- ▶ Assume the labels are encoded as $y \in \{-1, 1\}$.
- ▶ Another **problem**: $H(\vec{x})$ can be any real number.
 - ▶ Output is not necessarily -1 or 1.
- ▶ We need to turn output of $H(\vec{x})$ into -1 or 1.

Sign Function

- **Idea:** use the **sign function**.

$$\text{sign}(z) = \begin{cases} 1 & \text{if } z > 0 \\ -1 & \text{if } z < 0 \\ 0 & \text{if } z = 0 \end{cases}$$

Linear Classifiers

- ▶ We will still use linear predictors.
 - ▶ $H(\vec{x}; \vec{w}) = \text{Aug}(\vec{x}) \cdot \vec{w}$.
- ▶ But our final **predicted label** will be $\text{sign}(H(\vec{x}; \vec{w}))$.
 - ▶ If $H(\vec{x}) = 0$, predict either 1 or -1 (it's arbitrary).
- ▶ $\text{sign}(H(\vec{x}; \vec{w}))$ is called a **linear classifier**.
 - ▶ Takes in a feature vector and outputs a discrete label.
 - ▶ Sometimes called a **linear decision function**.


Interpretation: Weighted Vote

- ▶ A linear classifier is like a **weighted vote**.
- ▶ Each term $w_i x_i$ “votes” on the label.

$$H(\vec{X}) = w_0 + w_1 x_1 + w_2 x_2 + \dots + w_d x_d$$

- ▶ If the sum is:
 - ▶ positive: predict 1.
 - ▶ negative: predict -1.
 - ▶ zero: toss a coin, it's arbitrary!

The Prediction Surface




- ▶ $H(\vec{x})$ is a (hyper) plane.
- ▶ The place where $H(\vec{x}) = 0$ is the **decision boundary**.
- ▶ On one side, we predict 1.
- ▶ On the other, we predict -1.

Decision Boundary

- ▶ The **decision boundary** is the place where the output of $H(\vec{x})$ switches from “yes” to “no”.
- ▶ If H is a linear predictor and¹
 - ▶ $\vec{x} \in \mathbb{R}^1$, then the decision boundary is just a number.
 - ▶ $\vec{x} \in \mathbb{R}^2$, the boundary is a straight line.
 - ▶ $\vec{x} \in \mathbb{R}^d$, the boundary is a $d - 1$ dimensional (hyper) plane.

¹when plotted in the original feature coordinate space!

Magnitude of H




- The magnitude of $H(\vec{x})$ is **proportional** to the distance from the decision boundary.


Exercise

True or False: it's possible for two different linear prediction functions $H_1(\vec{x})$ and $H_2(\vec{x})$ to have the exact same decision boundary.


True



True



True



Another Useful Fact

- ▶ \vec{w} controls the orientation of the decision boundary.
 - ▶ A different \vec{w} gives a different decision boundary.
- ▶ Let $\vec{w}' = (w_1, \dots, w_d)$.
 - ▶ In other words, it is \vec{w} without the bias term w_0 .
- ▶ **Fact:** the decision boundary of $H(\vec{x}) = \text{Aug}(\vec{x}) \cdot \vec{w}$ is orthogonal to \vec{w}' .

Finding a Linear Classifier

- ▶ How do we find a good linear classifier?

ERM for Classification

- ▶ Step 1: choose a **hypothesis class**
 - ▶ We've chosen linear classifiers, $\text{sign}(\text{Aug}(\vec{x}) \cdot \vec{w})$.
- ▶ Step 2: choose a **loss function**
- ▶ Step 3: find H minimizing **empirical risk**
 - ▶ In case of linear predictors, equivalent to finding \vec{w} .

A First Idea

- ▶ Let's try using the same, familiar **square loss**.

DSC 140A

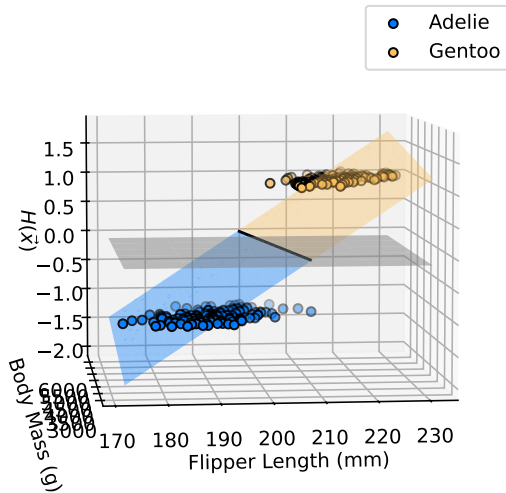
Probabilistic Modeling & Machine Learning

Lecture 6 | Part 2

Least Squares Classifiers

Classification as Regression

- ▶ We can think of classification as a special case of regression where the labels are always 1 or -1.
- ▶ **Goal:** find a prediction function $H(\vec{x})$ whose output is:
 - ▶ close to 1 for points from positive class.
 - ▶ close to -1 for points from negative class.



Least Squares Classifier

- ▶ **Idea:** least squares regression can be used for classification, too.
- ▶ The resulting algorithm is called the **least squares classifier**.

Linear Least Squares Classification

► To train:

- Given training data $(\vec{x}_1, y_1), \dots, (\vec{x}_n, y_n)$, with $y_i \in \{-1, 1\}$.
 1. Construct $n \times (d + 1)$ augmented design matrix, X .
 2. Solve the **normal equations**: $\vec{w}^* = (X^T X)^{-1} X^T \vec{y}$.

► To predict:

- Given a new point \vec{x} , predict $\text{sign}(\text{Aug}(\vec{x}) \cdot \vec{w}^*)$.


Square Loss for Classification

- ▶ We designed square loss for **regression**
- ▶ We can use it for **classification**.
- ▶ But it might not be the best choice.

Exercise

What is the **total** square loss of the predictor on the data?

Assume ● is class -1 and ● is class 1.



Observation

- ▶ The square loss penalizes points that are far from the decision boundary.
- ▶ **Even if they are correctly classified!**

Least Squares and Outliers




Figure 4.4 The left plot shows data from two classes, denoted by red crosses and blue circles, together with the decision boundary found by least squares (magenta curve) and also by the logistic regression model (green curve), which is discussed later in Section 4.3.2. The right-hand plot shows the corresponding results obtained when extra data points are added at the bottom left of the diagram, showing that least squares is highly sensitive to outliers, unlike logistic regression.

2

Another Loss?

- ▶ Least squares classifiers can work well in practice.
 - ▶ **Easy to implement!**
- ▶ But maybe a loss designed for classification will work better.

DSC 140A

Probabilistic Modeling & Machine Learning

Lecture 6 | Part 3

0-1 Loss

Empirical Risk Minimization

- ▶ Step 1: choose a **hypothesis class**
 - ▶ Let's assume we've chosen linear predictors
- ▶ Step 2: choose a **loss function**
- ▶ Step 3: minimize **expected loss (empirical risk)**

Another Loss Function

- ▶ What about the **0-1 loss**?
 - ▶ Loss = 0 if prediction is **correct**.
 - ▶ Loss = 1 if prediction is **incorrect**.
- ▶ More formally:

$$\ell_{0-1}(H(\vec{x}^{(i)}), y_i) = \begin{cases} 0 & \text{if } \text{sign}(H(\vec{x}^{(i)})) = y_i \\ 1 & \text{if } \text{sign}(H(\vec{x}^{(i)})) \neq y_i \end{cases}$$

Expected 0-1 Loss

- The expected 0-1 loss (empirical risk) has a nice interpretation:

$$R_{0-1}(H) = \frac{1}{n} \sum_{i=1}^n \begin{cases} 0 & \text{if } \text{sign}(H(\vec{x}^{(i)})) = y_i \\ 1 & \text{if } \text{sign}(H(\vec{x}^{(i)})) \neq y_i \end{cases}$$

Exercise

What is it?

Answer

- ▶ The empirical risk with respect to the 0-1 loss is the **misclassification rate** of the classifier.
 - ▶ That is, (1 - the **accuracy**)


$$R_{0-1}(H) = \frac{1}{n} \sum_{i=1}^n \begin{cases} 0 & \text{if } \text{sign}(H(\vec{x}^{(i)})) = y_i \\ 1 & \text{if } \text{sign}(H(\vec{x}^{(i)})) \neq y_i \end{cases}$$
$$= \frac{\# \text{ of } \textbf{incorrect} \text{ predictions}}{n}$$

ERM for the 0-1 Loss

- ▶ Minimizing the empirical risk with respect to the 0-1 loss is equivalent to **maximizing the accuracy**.
- ▶ That's exactly what we want!
- ▶ But there's a **problem**...

Exercise

What is the **gradient** of $R_{0-1}(H)$ with respect to the current \vec{w} ?



Answer

- ▶ The gradient of R_{0-1} is $\vec{0}$ almost everywhere.
- ▶ In other words, R_{0-1} is **flat** almost everywhere.
- ▶ This is a **problem** because **gradient descent** needs slope information to make progress.

Computationally Difficult

- ▶ It is **not feasible** to minimize 0-1 risk in general.
- ▶ More formally: **NP-Hard** to optimize expected 0-1 loss in general.³

³It is efficiently doable if the classes are linearly separable by finding convex hulls of each class. If non-separable, it is difficult.

Main Idea

It is computationally difficult (NP-Hard) to find a linear classifier with maximum accuracy, in general.

DSC 140A

Probabilistic Modeling & Machine Learning

Lecture 6 | Part 4

Perceptron Loss

Surrogate Loss

- ▶ We'd like to use the 0-1 loss, but it's not feasible.
- ▶ Instead, we use a **surrogate loss**.
- ▶ That is, a loss that is similar in spirit, but leads to easier optimization problems.

A New Loss

- ▶ No penalty if point is **correctly classified**.
 - ▶ Like the 0-1 loss.
- ▶ A penalty that grows with distance to decision boundary if point is **incorrectly classified**.
 - ▶ Unlike the 0-1 loss.
 - ▶ This will give us a non-zero gradient.

Perceptron Loss

- We call this loss the **perceptron loss**.


$$\ell_{\text{tron}}(H(\vec{x}), y) = \begin{cases} 0, & \text{sign}(H(\vec{x})) = y \\ |H(\vec{x})|, & \text{sign}(H(\vec{x})) \neq y \end{cases}$$

- Remember, $|H(\vec{x})|$ is **proportional** to distance from decision boundary.

Exercise

What is the **total** perceptron loss of the predictor on the data?

Assume ● is class -1 and ● is class 1.



Convexity?

- ▶ Is the perceptron loss **convex** in \vec{w} ?
- ▶ Trick:

$$\begin{aligned}\ell_{\text{tron}}(\text{Aug}(\vec{x}) \cdot \vec{w}, y) &= \begin{cases} 0, & \text{sign}(\text{Aug}(\vec{x}) \cdot \vec{w}) = y \\ |\text{Aug}(\vec{x}) \cdot \vec{w}|, & \text{sign}(\text{Aug}(\vec{x}) \cdot \vec{w}) \neq y \end{cases} \\ &= \max(0, -y \text{Aug}(\vec{x}) \cdot \vec{w})\end{aligned}$$

- ▶ **Fact:** Max of convex functions **is convex**.

ERM for the Perceptron

- **Goal:** minimize empirical risk w.r.t. perceptron loss for a linear predictor $H(\vec{x}) = \text{Aug}(\vec{x}) \cdot \vec{w}$.

$$\begin{aligned} R_{\text{tron}}(\vec{w}) &= \frac{1}{n} \sum_{i=1}^n \ell_{\text{tron}}(H(\vec{x}^{(i)}), y_i) \\ &= \frac{1}{n} \sum_{i=1}^n \begin{cases} 0, & \text{sign}(\text{Aug}(\vec{x}^{(i)}) \cdot \vec{w}) = y_i \\ |\text{Aug}(\vec{x}^{(i)}) \cdot \vec{w}|, & \text{sign}(\text{Aug}(\vec{x}^{(i)}) \cdot \vec{w}) \neq y_i \end{cases} \end{aligned}$$


Minimizing Perceptron Risk

- ▶ R_{tron} is **not differentiable**.
 - ▶ Because of the absolute value.
- ▶ But it is **convex**.
 - ▶ Since ℓ_{tron} is convex.
- ▶ We can minimize using **subgradient descent**.

A Subgradient of the Loss

- We need a subgradient of ℓ_{tron} .

$$\ell_{\text{tron}}(\text{Aug}(\vec{x}) \cdot \vec{w}, y) = \max(0, -y \text{Aug}(\vec{x}) \cdot \vec{w})$$



A Subgradient of the Loss

- ▶ We need a subgradient of ℓ_{tron} .

$$\ell_{\text{tron}}(\text{Aug}(\vec{x}) \cdot \vec{w}, y) = \max(0, -y \text{Aug}(\vec{x}) \cdot \vec{w})$$

- ▶ If $-y \text{Aug}(\vec{x}) \cdot \vec{w} > 0$, the gradient is $-y \text{Aug}(\vec{x})$.
- ▶ If $-y \text{Aug}(\vec{x}) \cdot \vec{w} < 0$, the gradient is $\vec{0}$.
- ▶ **Claim:** at $-y \text{Aug}(\vec{x}) \cdot \vec{w} = 0$, $\vec{0}$ is a subgradient.

Subgradient of the Loss

- We've found:

$$\begin{aligned} & \text{subgrad } \ell_{\text{tron}}(\text{Aug}(\vec{x}) \cdot \vec{w}, y) \\ &= \begin{cases} \vec{0}, & \text{if } -y \text{Aug}(\vec{x}) \cdot \vec{w} < 0 \\ -y \text{Aug}(\vec{x}), & \text{if } -y \text{Aug}(\vec{x}) \cdot \vec{w} > 0 \end{cases} \end{aligned}$$

- Or, equivalently:

$$\begin{aligned} & \text{subgrad } \ell_{\text{tron}}(\text{Aug}(\vec{x}) \cdot \vec{w}, y) \\ &= \begin{cases} \vec{0}, & \text{if } \text{sign}(\text{Aug}(\vec{x}) \cdot \vec{w}) = y \\ -y \text{Aug}(\vec{x}), & \text{if } \text{sign}(\text{Aug}(\vec{x}) \cdot \vec{w}) \neq y \end{cases} \end{aligned}$$

Subgradient of the Risk

- ▶ A subgradient of the risk is then:

$$\text{subgrad } R_{\text{tron}}(\vec{w}) =$$

$$\frac{1}{n} \sum_{i=1}^n \begin{cases} \vec{0}, & \text{sign}(\text{Aug}(\vec{x}^{(i)}) \cdot \vec{w}) = y_i \\ -y_i \text{Aug}(\vec{x}^{(i)}), & \text{sign}(\text{Aug}(\vec{x}^{(i)}) \cdot \vec{w}) \neq y_i \end{cases}$$

The Perceptron

► To train:

- Given training data $(\vec{x}_1, y_1), \dots, (\vec{x}_n, y_n)$, with $y_i \in \{-1, 1\}$.
- 1. Minimize $R_{\text{train}}(\vec{w})$ with, e.g., subgradient descent:

$$\vec{w}^{(t+1)} = \vec{w}^{(t)} - \eta(t) \times \frac{1}{n} \sum_{i=1}^n \begin{cases} \vec{0}, & \text{sign}(\text{Aug}(\vec{x}^{(i)}) \cdot \vec{w}) = y_i \\ -y_i \text{Aug}(\vec{x}^{(i)}), & \text{sign}(\text{Aug}(\vec{x}^{(i)}) \cdot \vec{w}) \neq y_i \end{cases}$$

► To predict:

- Given a new point \vec{x} , predict $\text{sign}(\text{Aug}(\vec{x}) \cdot \vec{w}^*)$.

DSC 140A

Probabilistic Modeling & Machine Learning


Lecture 6 | Part 5

Perceptron Demo: MNIST

Demo: MNIST

- ▶ MNIST is a classic machine learning data set.
- ▶ Many images of handwritten digits, 0-9.
- ▶ Multiclass classification problem.
- ▶ But we can make it binary: 3 vs. 7.

Example MNIST Digit



- ▶ Grayscale
- ▶ 28 x 28 pixels

MNIST Feature Vectors

- ▶ $28 \times 28 = 784$ pixels
- ▶ Each image is a vector in \mathbb{R}^{784}
- ▶ Each feature is intensity of single pixel
 - ▶ black $\rightarrow 0$, white $\rightarrow 255$
- ▶ A **very** simple representation.

Demo: MNIST

- ▶ Use only images of 3s and 7s.
- ▶ 4132 training images.
- ▶ 680 testing images.
- ▶ Some minor tuning.
 - ▶ Added random noise for robustness.
 - ▶ Picked classification threshold automatically.

Perceptron Learning

- ▶ Linear prediction function parameterized by \vec{w} .
- ▶ In this case, we can “reshape” \vec{w} to be same size as input image.

Weight Vector

- ▶ Recall that the prediction is a **weighted vote**:

$$H(\vec{x}) = \text{sign}(w_0 + w_1 x_1 + w_2 x_2 + \dots + w_{784} x_{784})$$


- ▶ Positive \rightarrow 7, Negative \rightarrow 3
- ▶ w_i is the weight of pixel i
 - ▶ positive: if this pixel is bright, I think this is a 7
 - ▶ negative: if this pixel is bright, I think this is a 3
 - ▶ magnitude: confidence in prediction

Perceptron Training




Perceptron Weight Vector


I predict that this is a 3!



I predict that this is a 7!



I predict that this is a 3!



Perceptron Results

- ▶ Test accuracy: 97.3%

Square Loss for Classification

- ▶ What if we use square loss for classification?
- ▶ We *can*, but will it work well?


Results: Least Squares

- ▶ Test Accuracy: 96.7% (marginally worse)


I think that this is a 3.



I think that this is a 7.




I think that this is a 7.



Results: Least Squares

- Misclassifications are telling.


I think that this is a 7.



I think that this is a 7.




I think that this is a 7.




Least Squares Weight Vector

- Can visualize weight of each pixel as an image.



Least Squares Weight Vector


I predict that this is a 7!



I predict that this is a 3!




I predict that this is a 7!




Some History

- ▶ Perceptrons were one of the first “machine learning” models.
- ▶ The basis of modern neural networks.

Rosenblatt's Perceptron





Next Time

- ▶ We solve linear classification, once and for all.