Linear and Logistic Regression

Dr. Xiaowei Huang

https://cgi.csc.liv.ac.uk/~xiaowei/

Up to now,

- Two Classical Machine Learning Algorithms
  - Decision tree learning
  - K-nearest neighbor
    - What is k-nearest-neighbor classification
    - How can we determine similarity/distance
    - Standardizing numeric features
    - Speeding up k-NN
      - edited nearest neighbour
      - k-d trees for nearest neighbour identification

Confidence for decision tree (example)

- Random forest:
  - multiple decision trees are trained, by using different resamples of your data.
  - Probabilities can be calculated by the proportion of decision trees which vote for each class.

- For example, if 8 out of 10 decision trees vote to classify an instance as positive, we say that the confidence is 8/10.

Here, the confidences of all classes add up to 1

Confidence for k-NN classification (example)

- Classification steps are the same, recall $\hat{y} \leftarrow \underset{v \in \text{values}(Y)}{\operatorname{argmax}} \sum_{i=1}^{k} \delta(v, y^{(i)})$
- Given a class $\hat{y}$, we compute

$$acc\_dist = \sum_{i=1}^{k} \delta(\hat{y}, y^{(i)}) \cdot distance(\hat{y}, y^{(i)})$$

Accumulated distance to the supportive instances

- apply sigmoid function on the reciprocal of the accumulated distance

$$confidence = \frac{1}{1 + e^{-\frac{1}{acc\_dist}}}$$

Here, the confidences of all classes may not add up to 1

Softmax?

Today’s Topics

- linear regression
- linear classification
- logistic regression

Recap: dot product in linear algebra

$$f_w(x) = w^T x$$

$$w = \begin{bmatrix} 2 \\ 3 \end{bmatrix} \qquad x = \begin{bmatrix} 1 \\ 4 \end{bmatrix}$$

$$w^T x = 2 * 1 + 3 * 4 = 14$$

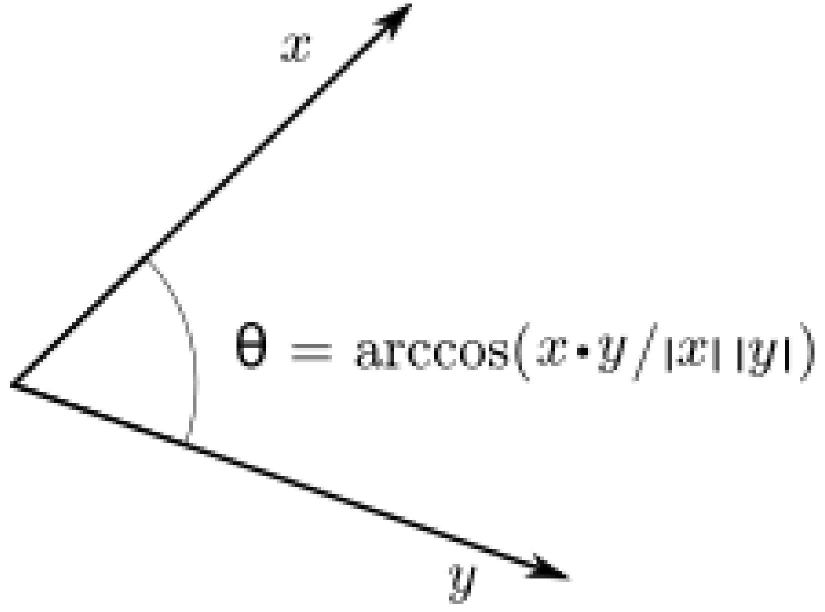

Geometric meaning: can be used to understand the angle between two vectors

Linear regression

Linear regression

- Given training data $\{(x^{(i)}, y^{(i)}) : 1 \leq i \leq m\}$ i.i.d. from distribution D

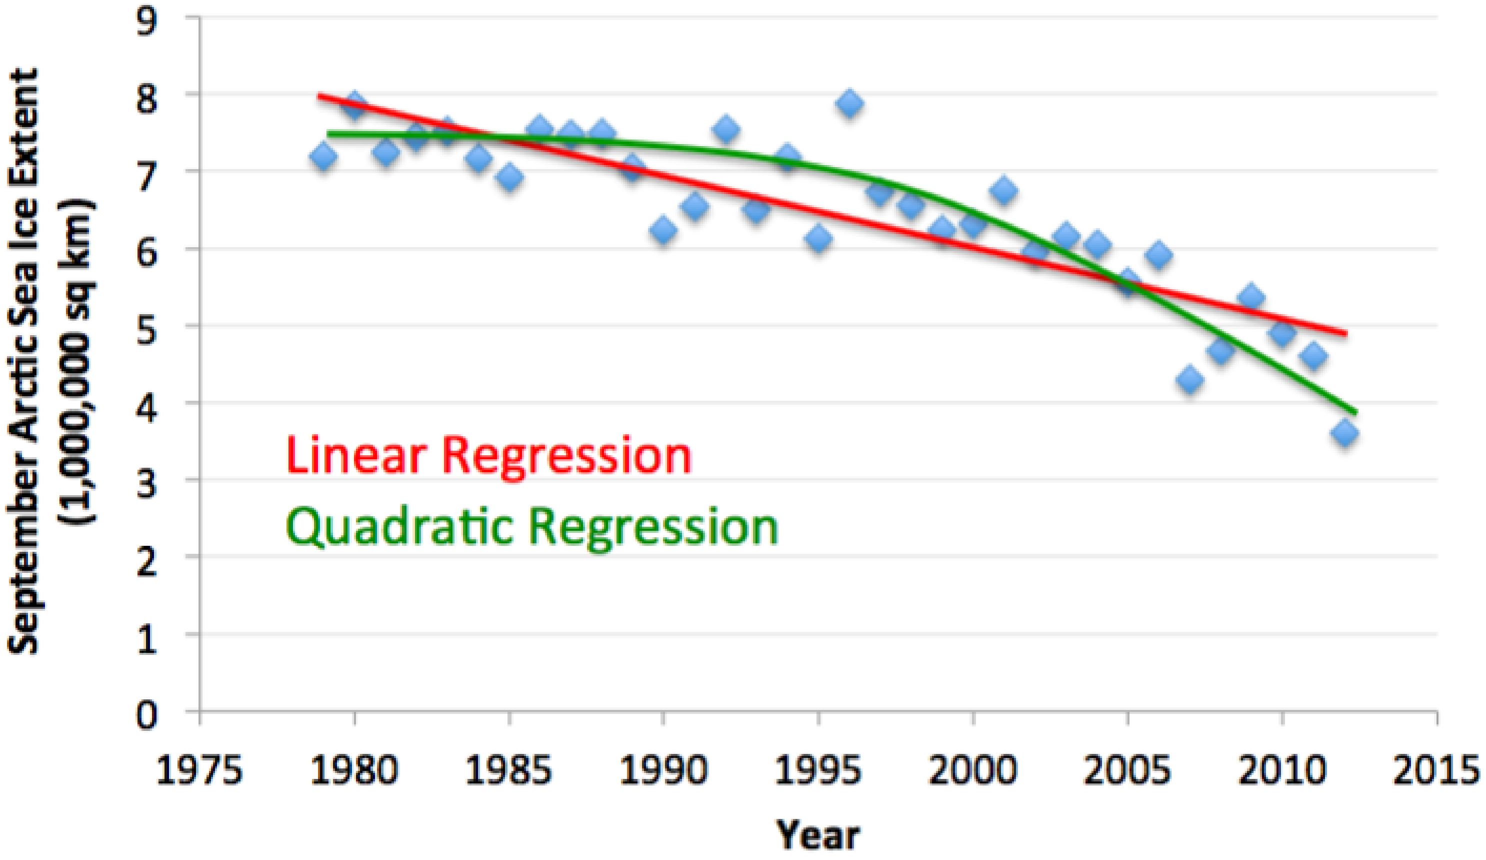

Recap: Consider the inductive bias of DT and k-NN learners

learner	hypothesis space bias	preference bias
ID3 decision tree	trees with single-feature, axis-parallel splits	small trees identified by greedy search
k-NN	Voronoi decomposition determined by nearest neighbors	instances in neighborhood belong to same class

Linear regression

- Given training data $\{(x^{(i)}, y^{(i)}) : 1 \leq i \leq m\}$ i.i.d. from distribution D
- Find $f_w(x) = w^T x$ that minimises

$$\hat{L}(f_w) = \frac{1}{m} \sum_{i=1}^{m} (w^T x^{(i)} - y^{(i)})^2$$

Hypothesis Class H

L^2 loss, or mean square error

Linear regression

- Given training data $\{(x^{(i)}, y^{(i)}) : 1 \leq i \leq m\}$ i.i.d. from distribution D
- Find $f_w(x) = w^T x$ that minimises

$$\hat{L}(f_w) = \frac{1}{m} \sum_{i=1}^{m} (w^T x^{(i)} - y^{(i)})^2$$

- where
  - $w^T x^{(i)} - y^{(i)}$ represents the error of instance $x^{(i)}$
  - $\sum_{i=1}^{m} (w^T x^{(i)} - y^{(i)})^2$ represents the square error of all training instances

So, $\frac{1}{m} \sum_{i=1}^{m} (w^T x^{(i)} - y^{(i)})^2$ represents the mean square error of all training instances

Linear regression

- Given training data $\{(x^{(i)}, y^{(i)}) : 1 \leq i \leq m\}$ i.i.d. from distribution D
- Find $f_w(x) = w^T x$ that minimises $\hat{L}(f_w) = \frac{1}{m}\sum_{i=1}^{m}(w^T x^{(i)} - y^{(i)})^2$

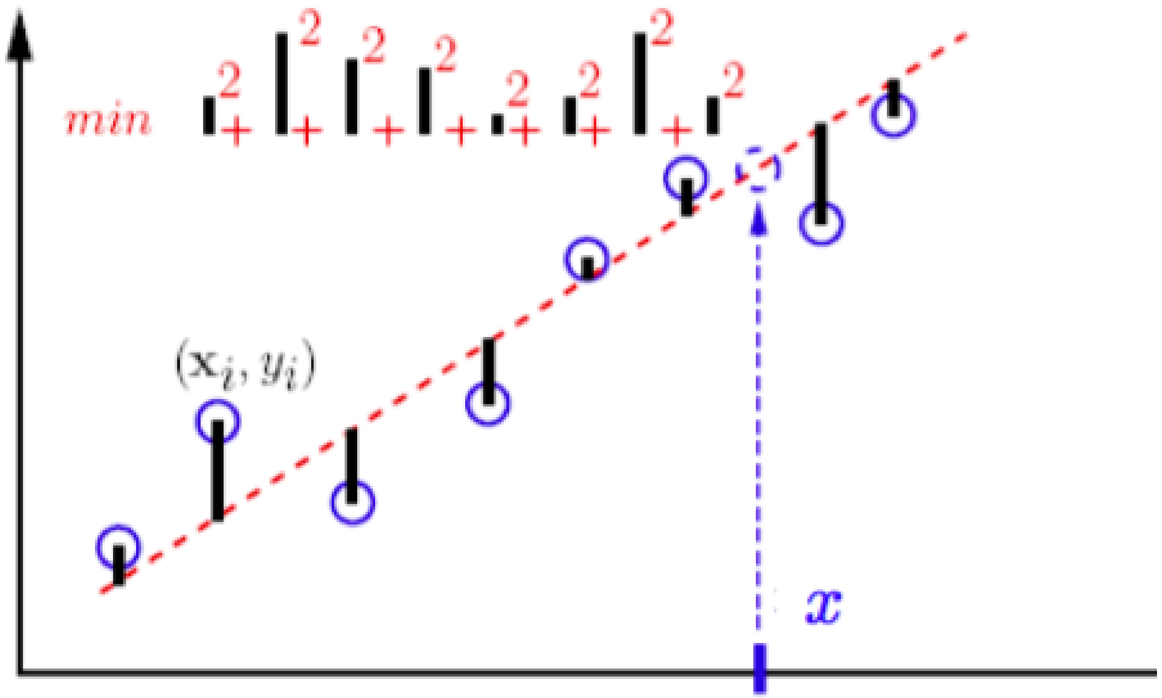

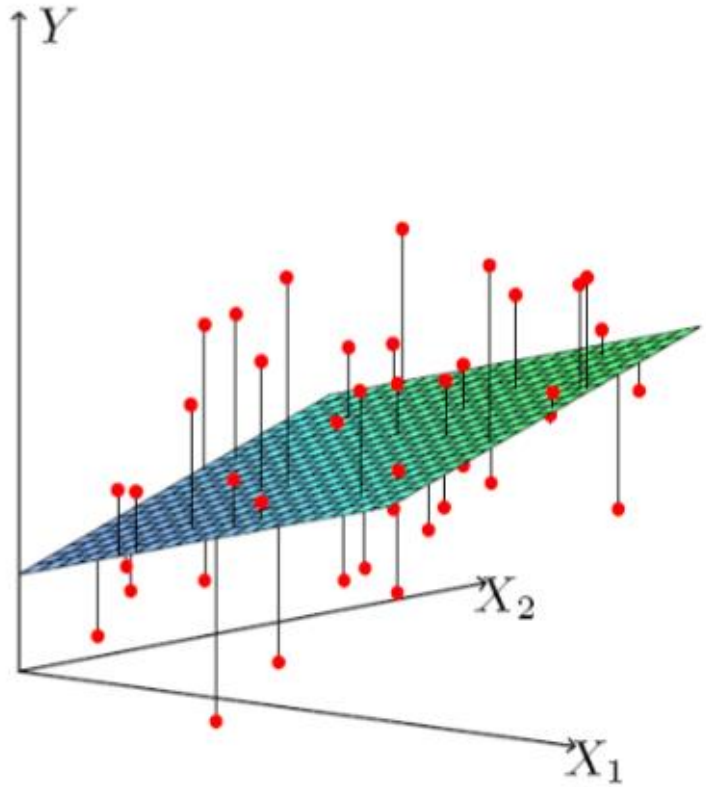

Organise feature data into matrix

- Let X be a matrix whose i-th row is $(x^{(i)})^T$

$$x^{(1)} = \begin{bmatrix} 182 \\ 87 \\ 11.3 \end{bmatrix} \quad x^{(2)} = \begin{bmatrix} 189 \\ 92 \\ 12.3 \end{bmatrix} \quad x^{(3)} = \begin{bmatrix} 178 \\ 79 \\ 10.6 \end{bmatrix} \quad x^{(4)} = \begin{bmatrix} 183 \\ 90 \\ 12.7 \end{bmatrix}$$

Football player example:
(height, weight, runningspeed)

$$X = \begin{bmatrix} 182 & 87 & 11.3 \\ 189 & 92 & 12.3 \\ 178 & 79 & 10.6 \\ 183 & 90 & 12.7 \end{bmatrix}$$

v1	**v2**	**v3**	**y**
182	87	11.3	325 (No)
189	92	12.3	344 (Yes)
178	79	10.6	350 (Yes)
183	90	12.7	320 (No)

Transform input matrix with weight vector

- Assume a function $f_w(x) = w^T x$ with weight vector $w = (1, -1, 20)$
  - Intuitively,
    - by 20, running speed is more important than the other two features, and
    - by -1, weight is negatively correlated to y

This is the parameter vector we want to learn

$$w^T x^{(1)} = \begin{bmatrix} 1 & -1 & 20 \end{bmatrix} * \begin{bmatrix} 182 \\ 87 \\ 11.3 \end{bmatrix} = 321$$

$$w^T x^{(2)} = \begin{bmatrix} 1 & -1 & 20 \end{bmatrix} * \begin{bmatrix} 189 \\ 92 \\ 12.3 \end{bmatrix} = 343.0$$

$$w^T x^{(3)} = 311 \qquad w^T x^{(4)} = 347$$

$$Xw = \begin{bmatrix} 321 \\ 343 \\ 311 \\ 347 \end{bmatrix}$$

Organise output into vector

- Let y be the vector $(y^{(1)}, ..., y^{(m)})^T$

v1	v2	v3	y
182	87	11.3	325
189	92	12.3	344
178	79	10.6	350
183	90	12.7	320

$$y = \begin{bmatrix} 325 \\ 344 \\ 350 \\ 320 \end{bmatrix}$$

Error representation

$$Xw = \begin{bmatrix} 321 \\ 343 \\ 311 \\ 347 \end{bmatrix} \qquad y = \begin{bmatrix} 325 \\ 344 \\ 350 \\ 320 \end{bmatrix}$$

- Square error of all instances

$$\sum_{i=1}^{m} (w^T x^{(i)} - y^{(i)})^2 \quad = ||Xw - y||_2^2$$

Linear regression : optimization

- Given training data $\{(x^{(i)}, y^{(i)}) : 1 \leq i \leq m\}$ i.i.d. from distribution D
- Find $f_w(x) = w^T x$ that minimises $\hat{L}(f_w) = \frac{1}{m}\sum_{i=1}^{m}(w^T x^{(i)} - y^{(i)})^2$

- Let X be a matrix whose i-th row is $(x^{(i)})^T$, y be the vector $(y^{(1)}, \ldots, y^{(m)})^T$

Now we knew where this comes from!

$$\hat{L}(f_w) = \frac{1}{m}\sum_{i=1}^{m}(w^T x^{(i)} - y^{(i)})^2 = \frac{1}{m}||Xw - y||_2^2$$

Solving this optimization problem will be introduced in later lectures.

Variant: Linear regression with bias

Linear regression with bias

- Given training data $\{(x^{(i)}, y^{(i)}) : 1 \le i \le m\}$ i.i.d. from distribution D
- Find $f_w(x) = w^T x + b$ that minimises the loss

Bias Term

- Reduce to the case without bias:
  - Let $w' = [w; b], x' = [x; 1]$
  - Then $f_{w,b}(x) = w^T x + b = (w')^T (x')$

Intuitively, every instance is extended with one more feature whose value is always 1, and we already know the weight for this feature, i.e., b

Linear regression with bias

- Think about bias $b = -330$ for the football player example

$$X'w' = \begin{bmatrix} -9 \\ 13 \\ -19 \\ 17 \end{bmatrix}$$

Can do a bit of exercise on this.

Variant: Linear regression with lasso penalty

Linear regression with lasso penalty

- Given training data $\{(x^{(i)}, y^{(i)}) : 1 \leq i \leq m\}$ i.i.d. from distribution D
- Find $f_w(x) = w^T x + b$ that minimises the loss

$$\hat{L}(f_w) = \frac{1}{m} \sum_{i=1}^{m} (w^T x^{(i)} - y^{(i)})^2 + \lambda |w|_1$$

lasso penalty: L^1 norm
of the parameter,
encourages sparsity

Variant: Evaluation Metrics

Evaluation Metrics

- Root mean squared error (RMSE)
- Mean absolute error (MAE) – average L^1 error
- R-square (R-squared)
- Historically all were computed on training data, and possibly adjusted after, but really should cross-validate

Linear classification

Linear classification

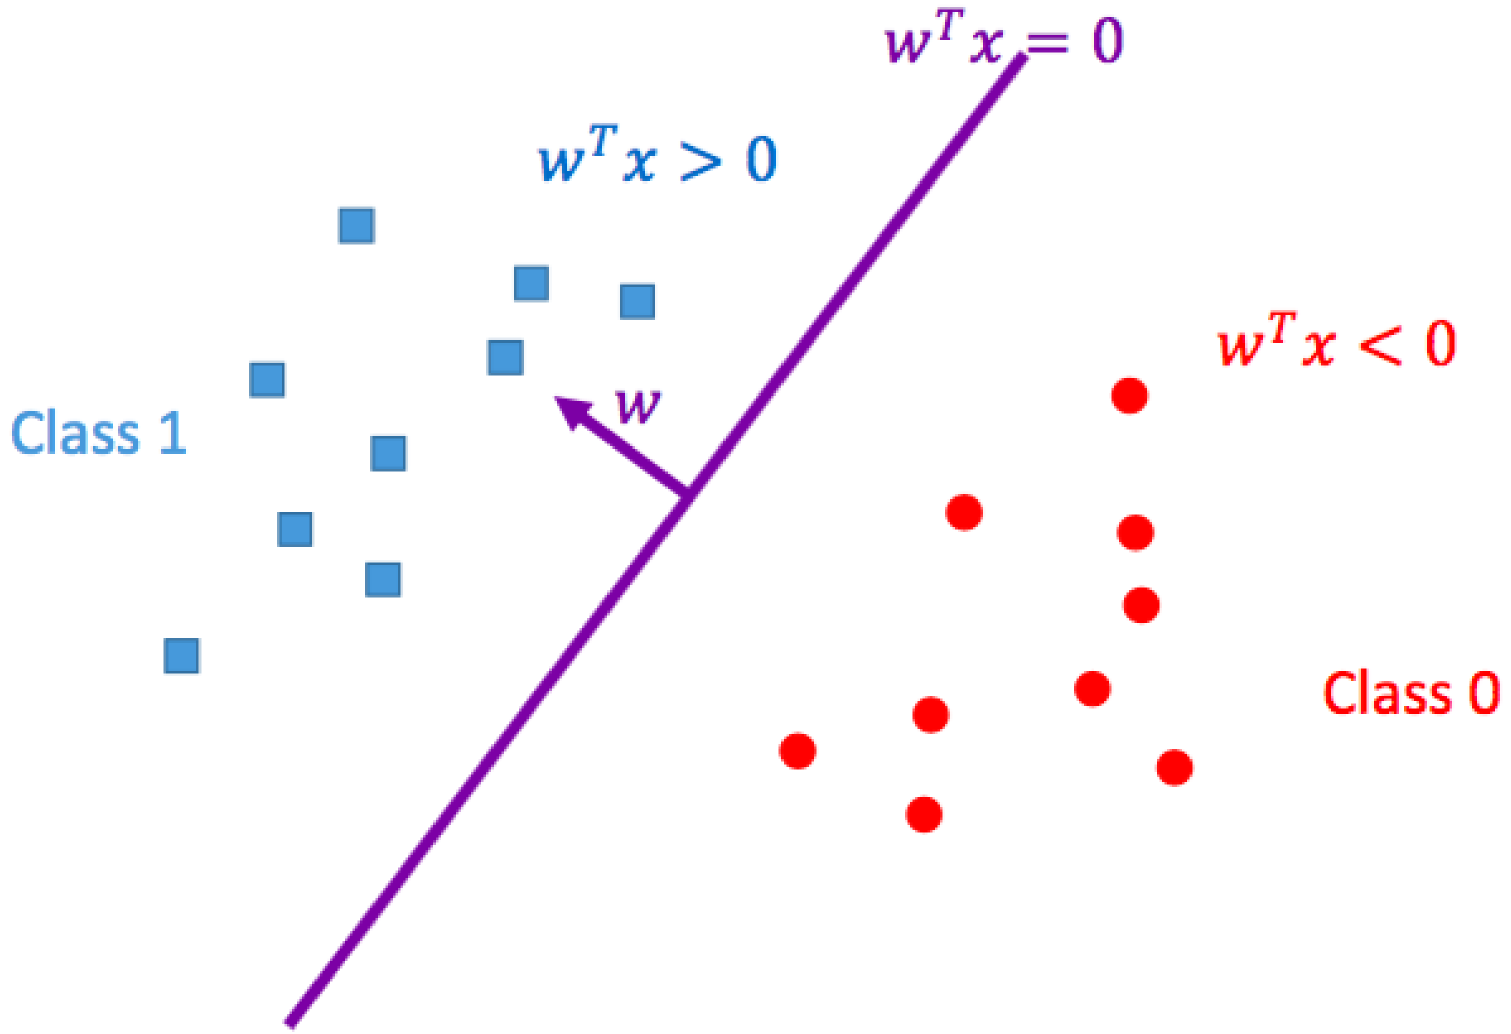

Linear classification: natural attempt

- Given training data $\{(x^{(i)}, y^{(i)}) : 1 \leq i \leq m\}$ i.i.d. from distribution D
- Hypothesis $f_w(x) = w^T x$
  - $y = 1$ if $w^T x > 0$
  - $y = 0$ if $w^T x < 0$

Piecewise Linear model $\mathcal{H}$

Still, w is the vector of parameters to be trained.

Or more formally, let $y = step(f_w(x)) = step(w^T x)$

where step(m) = 1, if m > 0 and
step(m) = 0, otherwise

But what is the optimization objective?

Linear classification: natural attempt

- Given training data $\{(x^{(i)}, y^{(i)}) : 1 \leq i \leq m\}$ i.i.d. from distribution D
- Find $f_w(x) = w^T x$ that minimises

$$\hat{L}(f_w) = \frac{1}{m} \sum_{i=1}^{m} \mathbb{I}[step(w^T x^{(i)}) \neq y^{(i)}]$$

- Drawback: difficult to optimize
  - NP-hard in the worst case

0-1 loss

loss = 0, i.e., no loss, when the classification is the same as its label.

loss =1, otherwise.

Linear classification

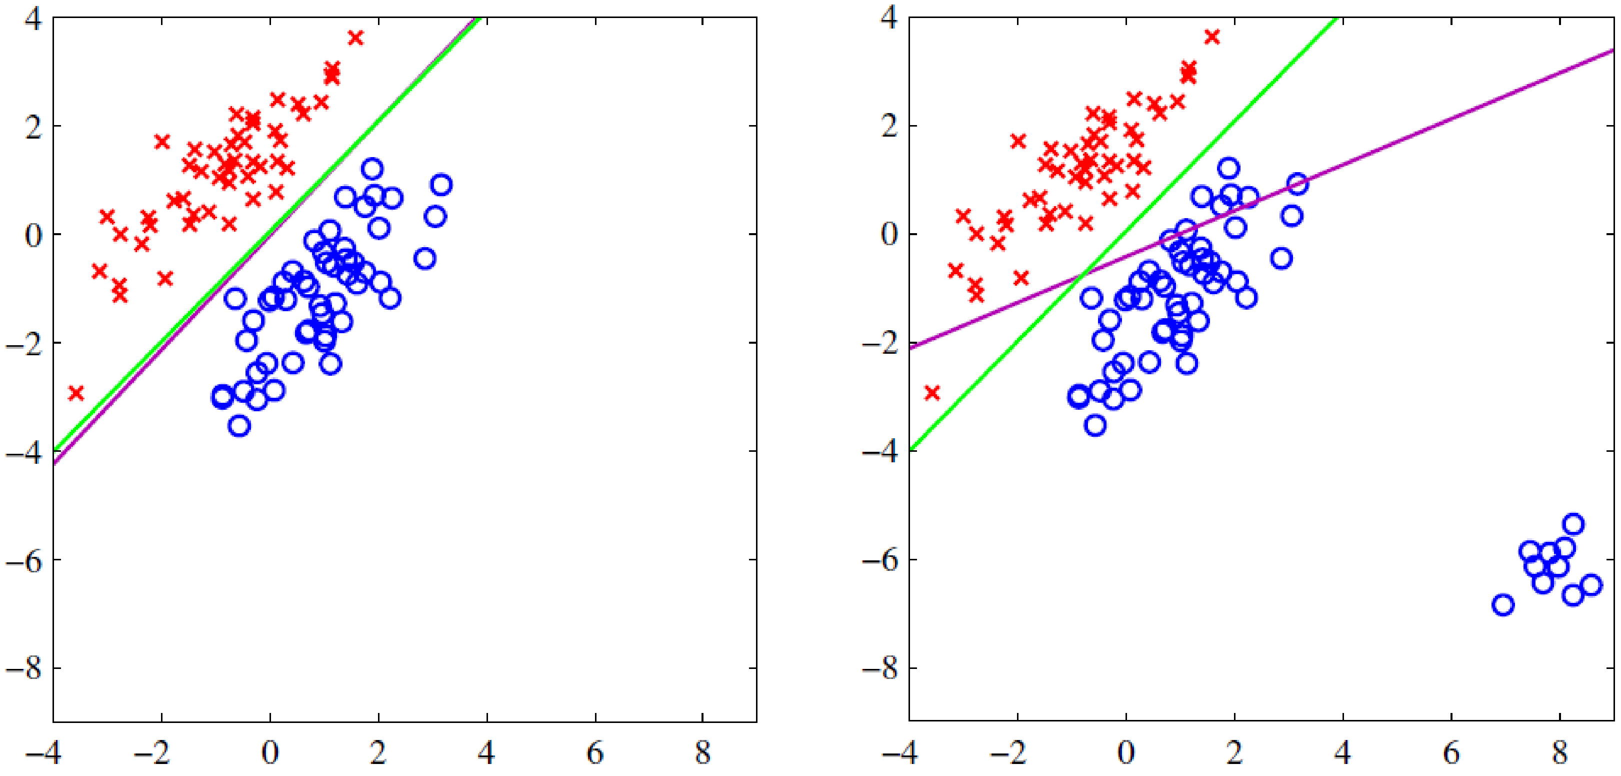

Figure 4.4 The left plot shows data from two classes, denoted by red crosses and blue circles, together with the decision boundary found by least squares (magenta curve) and also by the logistic regression model (green curve), which is discussed later in Section 4.3.2. The right-hand plot shows the corresponding 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.

Drawback: not robust to “outliers”

logistic regression

Why logistic regression?

- It's tempting to use the linear regression output as probabilities
- but it's a mistake because the output can be negative and greater than 1, whereas probability can not.

Logistic regression always outputs a value between 0 and 1

Compare the two

Linear regression

Linear classification

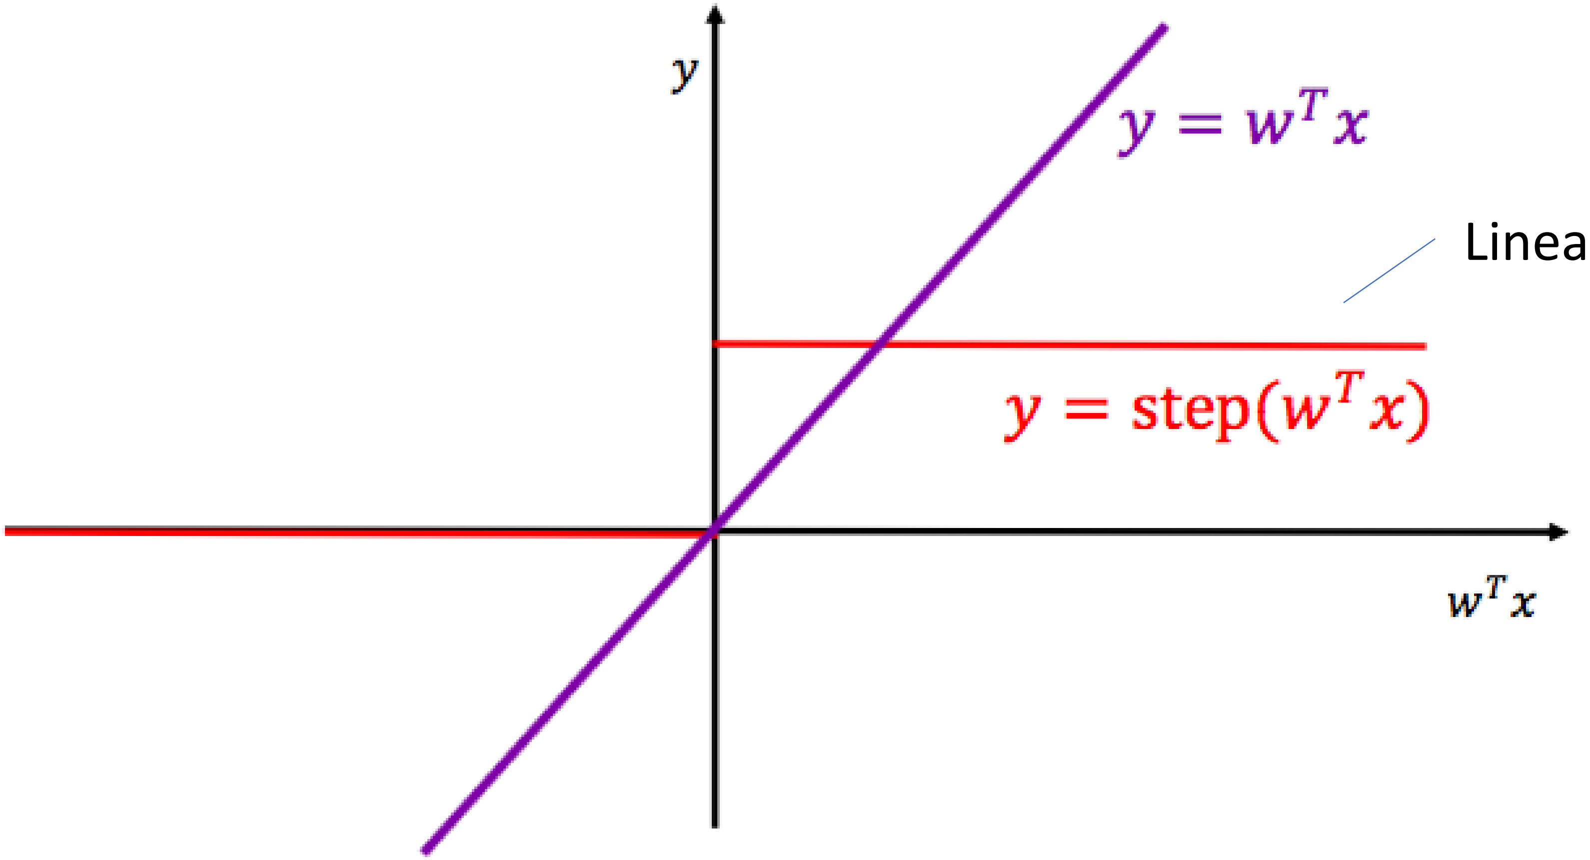

Between the two

- Prediction bounded in [0,1]
- Smooth

- Sigmoid:

$$\sigma(a) = \frac{1}{1+exp(-a)}$$

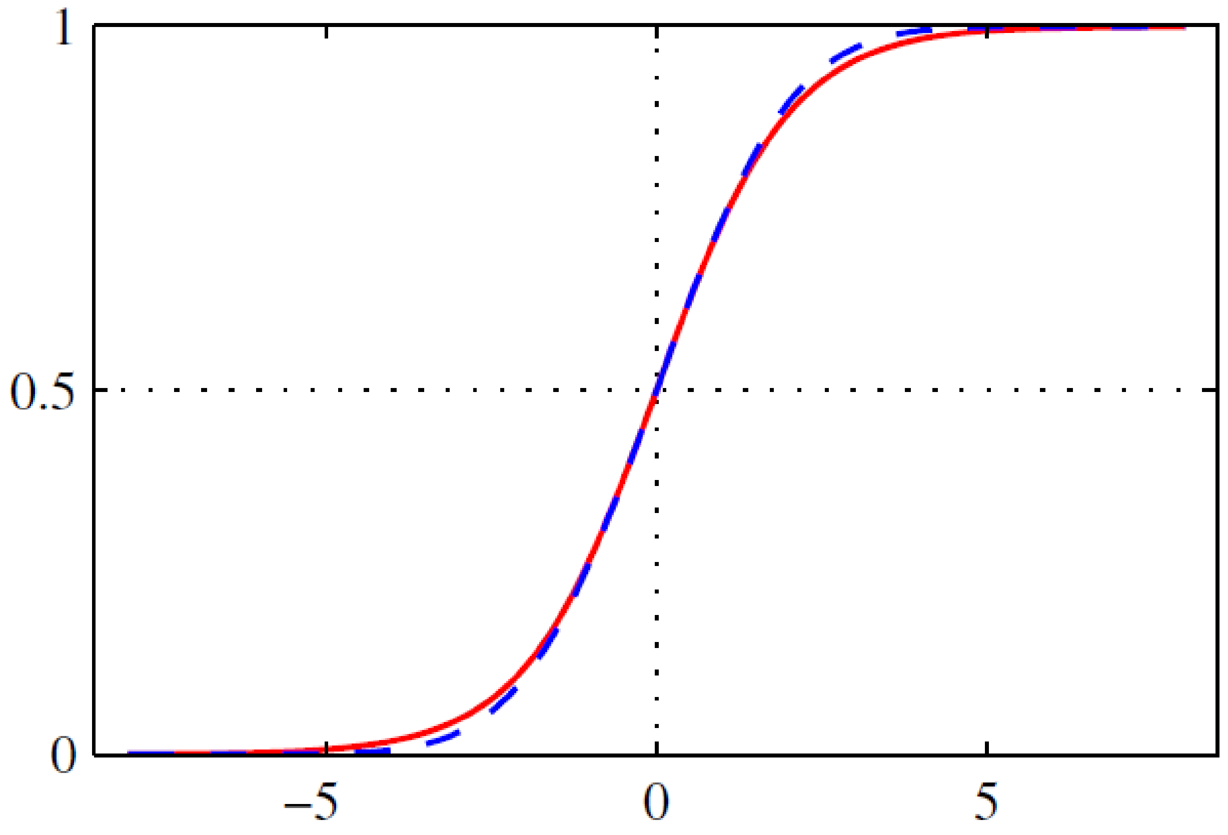

Linear regression: sigmoid prediction

- Squash the output of the linear function

$$Sigmoid(w^T x) = \sigma(w^T x) = \frac{1}{1+exp(-w^T x)}$$

New optimization objective

$$\hat{L}(f_w) = \frac{1}{m}\sum_{i=1}^{m}(\sigma(w^T x^{(i)}) - y^{(i)})^2$$

Question: Do we need to squash y?

Linear classification: logistic regression

$$\hat{L}(f_w) = \frac{1}{m}\sum_{i=1}^{m}(\sigma(w^T x^{(i)}) - y^{(i)})^2$$

- Is this the final?

We need a probability!

If $y^{(i)}$ is either 0 or 1, can we interpret $\sigma(w^T x^{(i)})$ as a probability value?

Linear classification: logistic regression

- A better approach: Interpret as a probability

Here we assume that y=0 or y=1

$$P_w(y = 1|x) = \sigma(w^T x)$$

$$P_w(y = 0 \mid x) = 1 - P_w(y = 1 \mid x) = 1 - \sigma(w^T x)$$

Conditional probability

Linear classification: logistic regression

- Find $f_w(x) = w^T x$ that minimises

$$\hat{L}(f_w) = \frac{1}{m}\sum_{i=1}^{m} \log P_w(y^{(i)}|x^{(i)})$$

Why log function used? To avoid numerical instability.

unfold

$$\hat{L}(f_w) = \frac{1}{m}\sum_{y^{(i)}=1} \log \sigma(w^T x^{(i)}) - \frac{1}{m}\sum_{y^{(i)}=0} \log[1-\sigma(w^T x^{(i)})]$$

Logistic regression:
MLE (maximum likelihood estimation) with sigmoid

Linear classification: logistic regression

- Given training data $\{(x^{(i)}, y^{(i)}) : 1 \leq i \leq m\}$ i.i.d. from distribution D
- Find w that minimises

$$\hat{L}(f_w) = \frac{1}{m} \sum_{y^{(i)}=1} \log \sigma(w^T x^{(i)}) - \frac{1}{m} \sum_{y^{(i)}=0} \log[1 - \sigma(w^T x^{(i)})]$$

No close form solution;
Need to use gradient descent

Why sigmoid function?

- Bounded

$$\sigma(a) = \frac{1}{1+\exp(-a)} \in (0,1)$$

- Symmetric

$$1 - \sigma(a) = \frac{\exp(-a)}{1+\exp(-a)} = \frac{1}{\exp(a)+1} = \sigma(-a)$$

- Gradient

$$\sigma'(a) = \frac{\exp(-a)}{(1+\exp(-a))^2} = \sigma(a)(1-\sigma(a))$$

Exercises

- Given the dataset and consider the mean square root error, if we have the following two linear functions:
  - $f_w(x) = 2x_1 + 1x_2 + 20x_3 - 330$
  - $f_w(x) = 1x_1 - 2x_2 + 23x_3 - 332$

  please answer the following questions:
  - (1) which model is better for linear regression?
  - (2) which model is better for linear classification by considering 0-1 loss for y^T=(No,Yes,Yes,No)?
  - (3) which model is better for logistic regression for y^T=(No,Yes,Yes,No)?
  - (4) According to the logistic regression of the first model, what is the prediction result of the first model on a new input (181,92,12.4).

x1	**x2**	**x3**	**y**
182	87	11.3	325
189	92	12.3	344
178	79	10.6	350
183	90	12.7	320