Supervised Learning

Given:

- a set of inputs features $X_1, \ldots, X_n$
- a set of target features $Y_1, \ldots, Y_k$
- a set of training examples where the values for the input features and the target features are given for each example
- a new example, where only the values for the input features are given

predict the values for the target features for the new example.

Supervised Learning

Given:

- a set of inputs features $X_1, \ldots, X_n$
- a set of target features $Y_1, \ldots, Y_k$
- a set of training examples where the values for the input features and the target features are given for each example
- a new example, where only the values for the input features are given

predict the values for the target features for the new example.

- classification when the Y_i are discrete
- regression when the Y_i are continuous

Example Data Representations

A travel agent wants to predict the preferred length of a trip, which can be from 1 to 6 days. (No input features).

Two representations of the same data:

— Y is the length of trip chosen.

— Each Y_i is an indicator variable that has value 1 if the chosen length is i, and is 0 otherwise.

Example	Y
e_1	1
e_2	6
e_3	6
e_4	2
e_5	1

Example	Y_1	Y_2	Y_3	Y_4	Y_5	Y_6
e_1	1	0	0	0	0	0
e_2	0	0	0	0	0	1
e_3	0	0	0	0	0	1
e_4	0	1	0	0	0	0
e_5	1	0	0	0	0	0

What is a prediction?

Evaluating Predictions

Suppose F is a feature and e is an example:

- val(e,F) is the value of feature F on example e.
- pval(e,F) is the predicted value of feature F on example e.
- The error of the prediction is a measure of how close $pval(e, Y)$ is to $val(e, Y)$.
- There are many possible errors that could be measured.

Measures of error

E is the set of examples. $\mathbf{T}$ is the set of target features.

- absolute error

$$\sum_{e \in E} \sum_{Y \in \mathbf{T}} |val(e, Y) - pval(e, Y)|$$

Measures of error

E is the set of examples. $\mathbf{T}$ is the set of target features.

- absolute error

$$\sum_{e \in E} \sum_{Y \in \mathbf{T}} |val(e, Y) - pval(e, Y)|$$

- sum of squares error

$$\sum_{e \in E} \sum_{Y \in \mathbf{T}} (val(e, Y) - pval(e, Y))^2$$

Measures of error

E is the set of examples. $\mathbf{T}$ is the set of target features.

- absolute error

$$\sum_{e \in E} \sum_{Y \in \mathbf{T}} \left| val(e, Y) - pval(e, Y) \right|$$

- sum of squares error

$$\sum_{e \in E} \sum_{Y \in \mathbf{T}} (val(e, Y) - pval(e, Y))^2$$

- worst-case error:

$$\max_{e \in E} \max_{Y \in \mathbf{T}} \left| val(e, Y) - pval(e, Y) \right| .$$

Measures of error

E is the set of examples. $\mathbf{T}$ is the set of target features.

- absolute error

$$\sum_{e \in E} \sum_{Y \in \mathbf{T}} |val(e, Y) - pval(e, Y)|$$

- sum of squares error

$$\sum_{e \in E} \sum_{Y \in \mathbf{T}} (val(e, Y) - pval(e, Y))^2$$

- worst-case error:

$$\max_{e \in E} \max_{Y \in \mathbf{T}} |val(e, Y) - pval(e, Y)| .$$

- A cost-based error takes into account costs of various errors.

Measures of error (cont.)

When target features are $\{0,1\}$:

- likelihood of the data

$$\prod_{e \in E} \prod_{Y \in \mathbf{T}} pval(e,Y)^{val(e,Y)} (1 - pval(e,Y))^{(1-val(e,Y))}$$

Measures of error (cont.)

When target features are $\{0, 1\}$:

- likelihood of the data

$$\prod_{e \in E} \prod_{Y \in \mathbf{T}} pval(e, Y)^{val(e,Y)} (1 - pval(e, Y))^{(1 - val(e,Y))}$$

- entropy (number of bits to encode the data given a code based on *pval*)

$$-\sum_{e \in E} \sum_{Y \in \mathbf{T}} [val(e, Y) \log pval(e, Y) + (1 - val(e, Y)) \log(1 - pval(e, Y))]$$

Information theory overview

- A **bit** is a binary digit.
- 1 bit can distinguish 2 items
- k bits can distinguish 2^k items
- n items can be distinguished using $\log_2 n$ bits
- Can you do better?

Information and Probability

Let's design a code to distinguish elements of $\{a, b, c, d\}$ with

$$P(a) = \frac{1}{2}, P(b) = \frac{1}{4}, P(c) = \frac{1}{8}, P(d) = \frac{1}{8}$$

Consider the code:

a	0	b	10	c	110	d	111

This code sometimes uses 1 bit and sometimes uses 3 bits. On average, it uses

$$\begin{aligned} & P(a) \times 1 + P(b) \times 2 + P(c) \times 3 + P(d) \times 3 \\ = \; & \frac{1}{2} + \frac{2}{4} + \frac{3}{8} + \frac{3}{8} = 1\frac{3}{4} \text{ bits.} \end{aligned}$$

The string *aacabbda* has code 00110010101110.

Information Content

- To identify x, you need $-\log_2 P(x)$ bits.
- If you have a distribution over a set and want to a identify a member, you need the expected number of bits:

$$\sum_x -P(x) \times \log_2 P(x).$$

  This is the information content or entropy of the distribution.
- The expected number of bits it takes to describe a distribution given evidence e:

$$I(e) = \sum_x -P(x|e) \times \log_2 P(x|e).$$

Information Gain

If you have a test that can distinguish the cases where α is true from the cases where α is false, the information gain from this test is:

$$I(\mathit{true}) - (P(\alpha) \times I(\alpha) + P(\neg\alpha) \times I(\neg\alpha)).$$

- $I(\mathit{true})$ is the expected number of bits needed before the test
- $P(\alpha) \times I(\alpha) + P(\neg\alpha) \times I(\neg\alpha)$ is the expected number of bits after the test.

Linear Predictions

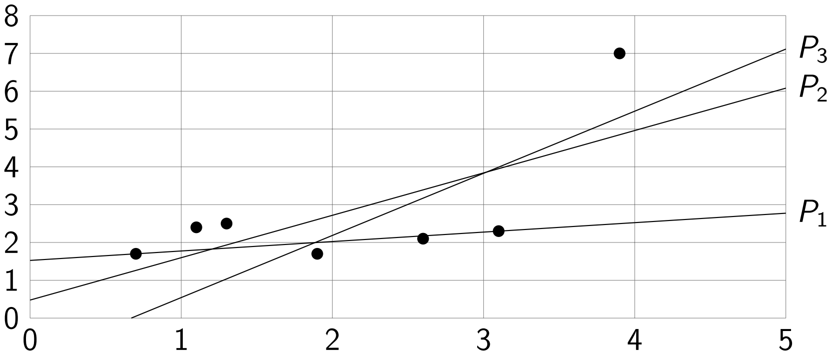

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is the mean (average) value of Y.

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is the mean (average) value of Y.
- The prediction that minimizes the absolute error on E is

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is the mean (average) value of Y.
- The prediction that minimizes the absolute error on E is the median value of Y.

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is the mean (average) value of Y.
- The prediction that minimizes the absolute error on E is the median value of Y.
- The prediction that minimizes the worst-case error on E is

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is the mean (average) value of Y.
- The prediction that minimizes the absolute error on E is the median value of Y.
- The prediction that minimizes the worst-case error on E is

$$(maximum + minimum)/2$$

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is the mean (average) value of Y.
- The prediction that minimizes the absolute error on E is the median value of Y.
- The prediction that minimizes the worst-case error on E is

$$(maximum + minimum)/2$$

- When Y has domain $\{0, 1\}$, the prediction that maximizes the likelihood on E is

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is the mean (average) value of Y.
- The prediction that minimizes the absolute error on E is the median value of Y.
- The prediction that minimizes the worst-case error on E is

$$(maximum + minimum)/2$$

- When Y has domain $\{0, 1\}$, the prediction that maximizes the likelihood on E is the empirical probability.

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is the mean (average) value of Y.
- The prediction that minimizes the absolute error on E is the median value of Y.
- The prediction that minimizes the worst-case error on E is

$$(maximum + minimum)/2$$

- When Y has domain $\{0, 1\}$, the prediction that maximizes the likelihood on E is the empirical probability.
- When Y has domain $\{0, 1\}$, the prediction that minimizes the entropy on E is

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is the mean (average) value of Y.
- The prediction that minimizes the absolute error on E is the median value of Y.
- The prediction that minimizes the worst-case error on E is

$$(maximum + minimum)/2$$

- When Y has domain $\{0,1\}$, the prediction that maximizes the likelihood on E is the empirical probability.
- When Y has domain $\{0,1\}$, the prediction that minimizes the entropy on E is the empirical probability.

Point Estimates

Suppose there is a single numerical feature, Y. Let E be the training examples. There is a single prediction for all examples.

- The prediction that minimizes the sum of squares error on E is the mean (average) value of Y.
- The prediction that minimizes the absolute error on E is the median value of Y.
- The prediction that minimizes the worst-case error on E is

$$(maximum + minimum)/2$$

- When Y has domain $\{0, 1\}$, the prediction that maximizes the likelihood on E is the empirical probability.
- When Y has domain $\{0, 1\}$, the prediction that minimizes the entropy on E is the empirical probability.

But that doesn't mean that these predictions minimize the error for future predictions.

Training and Test Sets

To evaluate how well a learner will work on future predictions, we divide the examples into:

- training examples that are used to train the learner
- test examples that are used to evaluate the learner

...these must be kept separate.

Learning Probabilities

- Empirical probabilities do not make good predictors when evaluated by likelihood or entropy.
- Why?

Learning Probabilities

- Empirical probabilities do not make good predictors when evaluated by likelihood or entropy.
- Why? A probability of zero means "impossible" and has infinite cost if there is one true case in test set.

Learning Probabilities

- Empirical probabilities do not make good predictors when evaluated by likelihood or entropy.
- Why? A probability of zero means "impossible" and has infinite cost if there is one true case in test set.
- Solution: add (non-negative) pseudo-counts to the data. Suppose n_i is the number of examples with $X = v_i$, and c_i is the pseudo-count:

$$P(X = v_i) = \frac{c_i + n_i}{\sum_{i'} c_{i'} + n_{i'}}$$

- Pseudo-counts convey prior knowledge. Consider: "how much more would I believe v_i if I had seen one example with v_i true than if I has seen no examples with v_i true?"