9/9 Lecture Notes

Admin: Change up seats, Lab 1 due tmr

Data repr. and featurization

Features (input) ex temp

Label (output) ex will you play outside.

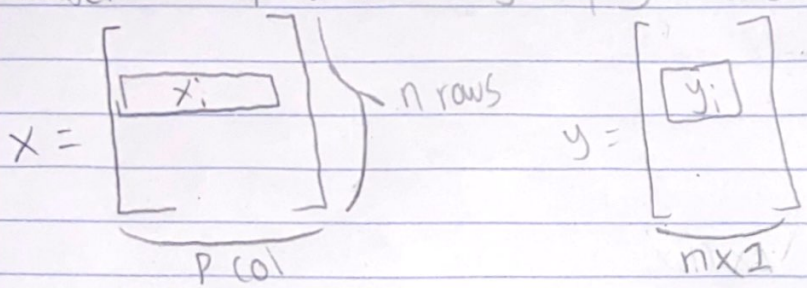

One value of y per data point, can be cont./discrete

Feature: name → shape

values → circle

vector → $x = [x_1, x_2, \cdots x_p]$

Featurization (make numerical) ex False → 0

↳ so comp. can interpret them True → 1

More ex. on slides Can map spectrum to values ex sunny → 1

rainy → 0

This process is done by Data Scientists

to suit the model best.

Model: distribution that captures data

ex decision tree (on slide)

Params	weather – feature		Data	weather	tennis
	sun /	\ rain – feat. value		S	Y
	Y	N – labels		R	N
	[1, 2]	[2, 0] – label counts		R	N
	2/3	2/2 – Accuracy		S	Y
		(80% overall)		S	N

Normal distribution model

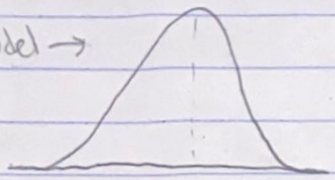

mean, Variance] Model Params

Linear models:

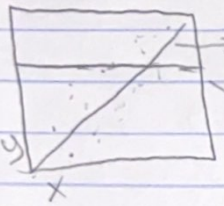

→ Model (slope = m) $y = mx + b$

prediction Based on model

$m, b \rightarrow$ model params

Handout thoughts: Q_2 - can featurize by binarizing (turning to 1s and 0s)

Q_4 - Model is "perfectly classified" depending on meaning

↳ if all given data points are classified

↳ if all possible data values are classified

Purpose of models:

Make predictions based on data

Visualize data

Help humans make choices based on data

↳ ex, can see that a certain feature aligns more with output

Linear models:

feature $\vec{x}$ eq: $h_{\vec{w}}(x) = w_0 + w_1 x = \hat{y}$ residual: $y_i - \hat{y}_i$

output y (prediction)

want to minimize RSS $\sum_{i=1} (y_i - \hat{y}_i)^2$

(error variance)

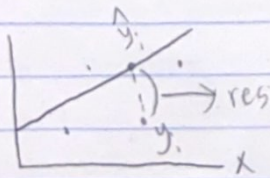

→ residual $(y_i - \hat{y}_i)$ → difference between prediction and data

***

Next lecture:

Use CL Args to input data

↳ parse_args package

Demo on website

Notes: can use — in arg name to make it optional

args are read in as string

Can set default arg value

Goals of fitting a linear model: 1) which features impact y

2) association between X and y

ex.

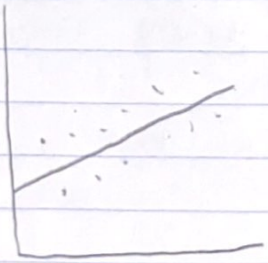

→ Model tells us as sales goes up, so does budget

Linear Regression

Used when output is continuous

Learned model: linear function mapping X to y (including feature weights and bias)

Goal: to minimize RSS/SSE (error)

model params

"Simple" lin. regress. → one feature

$h_{\vec{w}} = w_0 + w_1 x = \hat{y}$

weights

$\vec{w} = \begin{bmatrix} w_0 \\ w_1 \end{bmatrix}$

Linear model may not always be most accurate
↳ could use polynomial curve to better fit data
MSE (average error)

Performance may vary from training data to testing data
"Overfitting" - over performs on training data, underperforms on new data