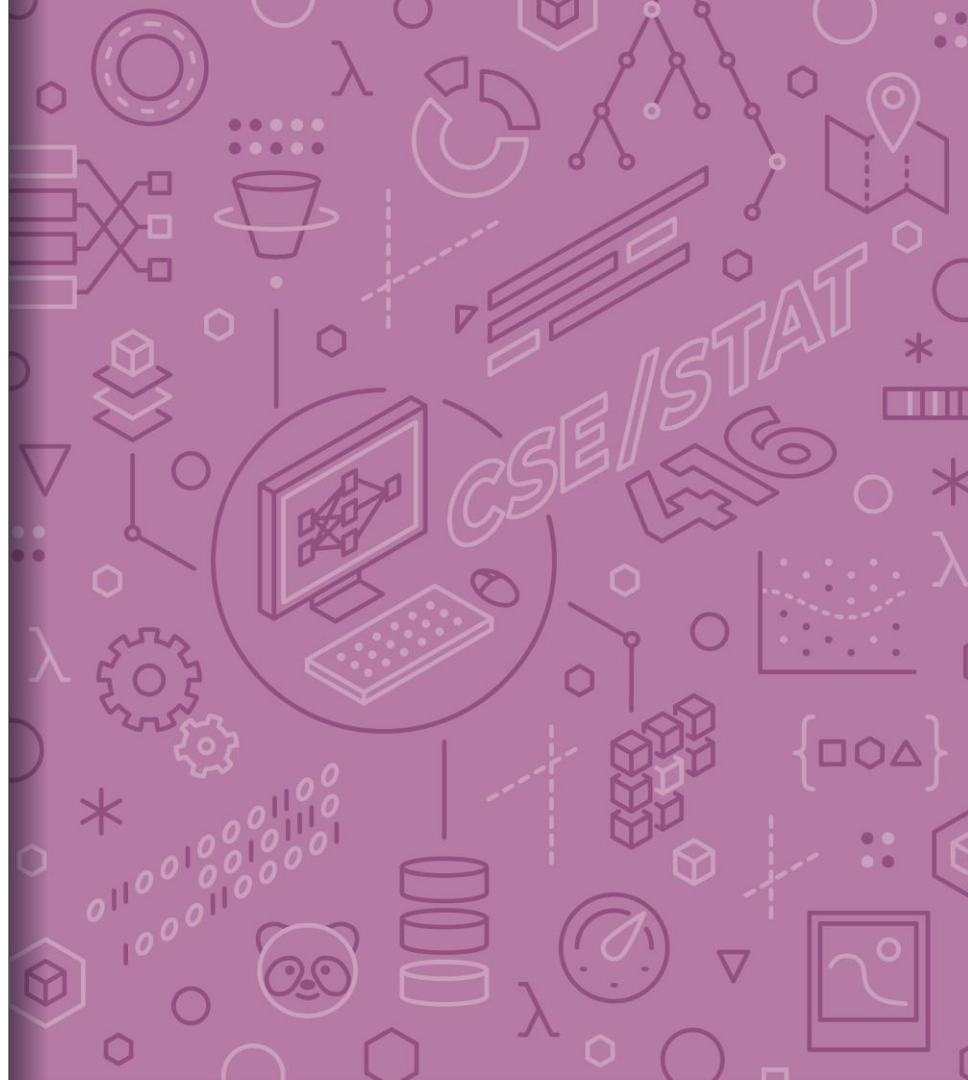


Logistic Regression

Adapted from Hunter Schafer's slides



Administrivia

- Week 5: Other ML models for classification
- Week 6: Deep Learning
- Homework 2 due yesterday
 - Up to Thurs 11:59PM with late days
- HW3 Released today, due Tues 7/19 11:59PM
- Next week's homework, **HW4, will allow groups of up to 2 for the programming part!**
 - See Ed for information about group formation!
- LR4 Due Fri 11:59PM



HW3

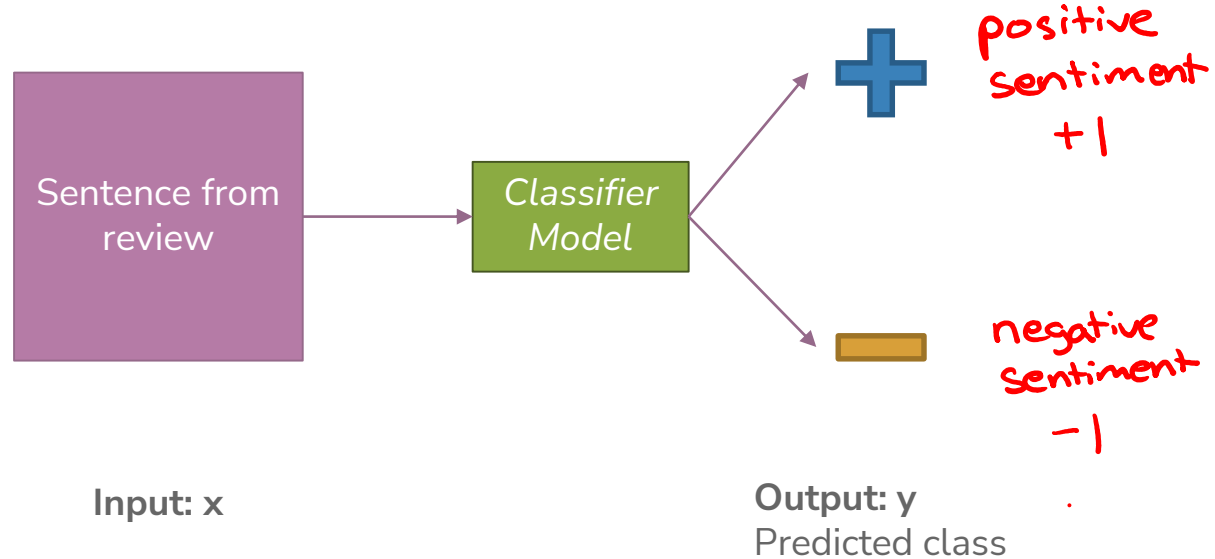
Walkthrough

Recap: Intro to Classification

*Continuing from
Lec5*

Sentiment Classifier

In our example, we want to classify a restaurant review as positive or negative.



Converting Text to Numbers (Vectorizing):

Bag of Words

- **Idea:** One feature per word!

Example: "Sushi was great, the food was awesome, but the service was terrible"

sushi	was	great	the	food	awesome	but	service	terrible
1	3	1	2	1	1	1	1	1

This **has** to be too simple, right?

- Stay tuned (today and Wed) for issues that arise and how to address them 😊

Attempt 3: Linear Classifier

(Another
View)

Idea: Only predict the sign of the output!

$$\text{Predicted Sentiment} = \hat{y} = \text{sign}(\text{Score}(x))$$

Linear Classifier

Input x : Sentence from review

- Compute $\text{Score}(x)$
- If $\text{Score}(x) > 0$: \leftarrow Threshold
 - $\hat{y} = +1$
- Else:
 - $\hat{y} = -1$

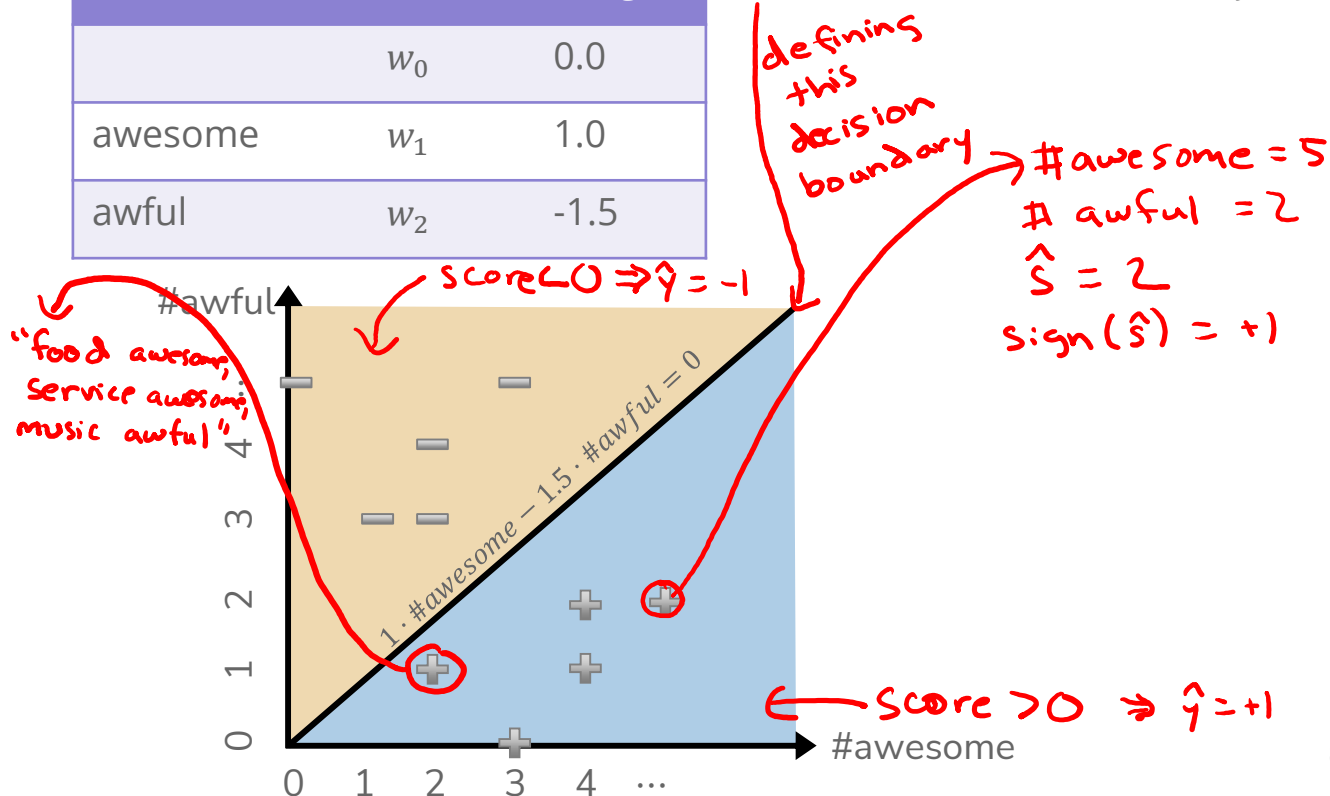
Earlier Example :
 $\text{Score}(x) = 2$
 $\hat{y} = +1$

Decision Boundary

Consider if only two words had non-zero coefficients

Word	Coefficient	Weight
	w_0	0.0
awesome	w_1	1.0
awful	w_2	-1.5

$$\hat{s} = 1 \cdot \#awesome - 1.5 \cdot \#awful$$



Classification Error

Ratio of examples where there was a mistaken prediction

What's a mistake?

- If the true label was positive ($y = +1$), but we predicted negative ($\hat{y} = -1$) \rightarrow False Negative
- If the true label was negative ($y = -1$), but we predicted positive ($\hat{y} = +1$) \rightarrow False Positive

Classification Error

$$\frac{\text{\# mistakes}}{\text{\# examples}} = \frac{\sum_{i=1}^n \mathbb{1}\{y_i \neq \hat{y}_i\}}{n}$$

Classification Accuracy

$$\frac{\text{\# correct}}{\text{\# examples}} = \frac{\sum_{i=1}^n \mathbb{1}\{y_i = \hat{y}_i\}}{n} = 1 - \text{error}$$

Confusion Matrix

For binary classification, there are only two types of mistakes

- $\hat{y} = +1, y = -1$
- $\hat{y} = -1, y = +1$

Generally we make a **confusion matrix** to understand mistakes.

Complete the sentence: "my prediction was a..."

		<u>Predicted Label</u>	
		+	-
<u>True Label</u>	+	True Positive (TP)	False Negative (FN)
	-	False Positive (FP)	True Negative (TN)

Tip on remembering: complete the sentence "My prediction was a ..."

Binary Classification Measures

Notation

- $C_{TP} = \#TP$, $C_{FP} = \#FP$, $C_{TN} = \#TN$, $C_{FN} = \#FN$
- $N = C_{TP} + C_{FP} + C_{TN} + C_{FN}$
- $N_P = C_{TP} + C_{FN}$, $N_N = C_{FP} + C_{TN}$

Error Rate

$$\frac{C_{FP} + C_{FN}}{N}$$

Accuracy Rate

$$\frac{C_{TP} + C_{TN}}{N}$$

False Positive rate (FPR)

$$\frac{C_{FP}}{N_N}$$

False Negative Rate (FNR)

$$\frac{C_{FN}}{N_P}$$

True Positive Rate or Recall

$$\frac{C_{TP}}{N_P}$$

Precision

$$\frac{C_{TP}}{C_{TP} + C_{FP}}$$

F1-Score

$$2 \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

[See more!](#)

Precision & Recall

Two particularly important metrics in binary classification are:

Precision: Of the ones I predicted positive, how many of them were actually positive?

- How precise is my model in its predictions?
- $TP / (TP + FP)$

Recall: Of all the things that are truly positive, how many of them did I correctly predict as positive?

- How good is your model at recalling the patterns in the training data?
- $TP / (TP + FN)$



Precision

What fraction of the examples I predicted positive were correct?

Sentences predicted to be positive:

$$\hat{y}_i = +1$$

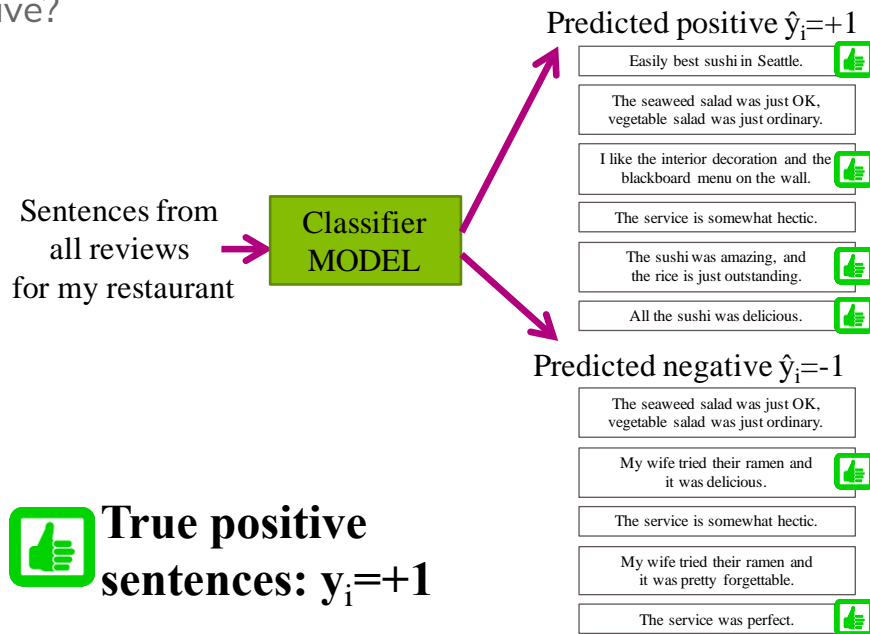
Easily best sushi in Seattle.	<input checked="" type="checkbox"/>
The seaweed salad was just OK, vegetable salad was just ordinary.	<input type="checkbox"/>
I like the interior decoration and the blackboard menu on the wall.	<input checked="" type="checkbox"/>
The service is somewhat hectic.	<input type="checkbox"/>
The sushi was amazing, and the rice is just outstanding.	<input checked="" type="checkbox"/>
All the sushi was delicious.	<input checked="" type="checkbox"/>

Only 4 out of 6
sentences
predicted to be
positive are
actually **positive**

$$precision = \frac{C_{TP}}{C_{TP} + C_{FP}}$$

Recall

Of the truly positive examples, how many were predicted positive?



$$recall = \frac{C_{TP}}{N_P} = \frac{C_{TP}}{C_{TP} + C_{FN}}$$

Precision & Recall

There is a tradeoff between precision and recall!

An optimistic model will predict almost everything as positive

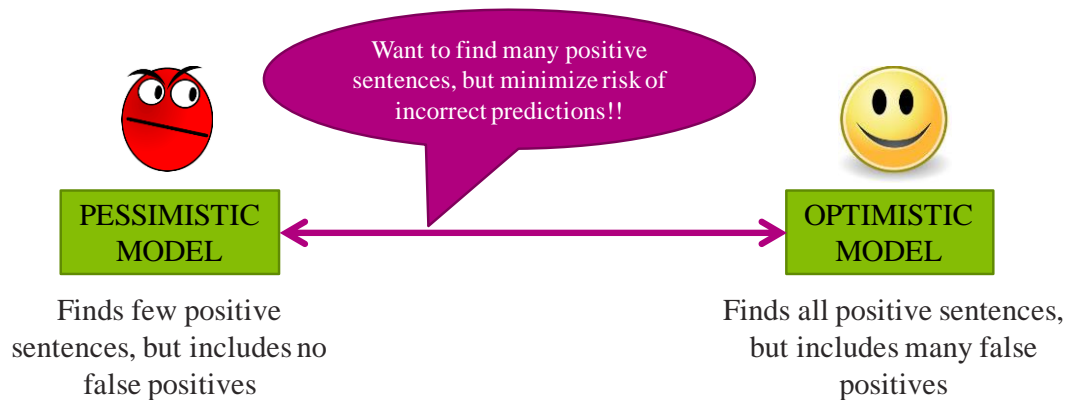


- High recall, low precision

A pessimistic model will predict almost everything as negative



- High precision, low recall



Multiclass Confusion Matrix

Consider predicting (*Healthy, Cold, Flu*)

		Predicted Label		
		Healthy	Cold	Flu
True Label	Healthy	60	8	2
	Cold	4	12	4
	Flu	0	2	8

Think

1 min

pollev.com/cs416

Suppose we trained a classifier and computed its confusion matrix on the training dataset. **Is there a class imbalance in the dataset and if so, which class has the highest representation?**

		Predicted Label		
		Pupper	Doggo	Woofers
True Label	Pupper	2	27	4
	Doggo	4	25	4
	Woofers	1	30	2

1:00

Think

2 min

pollev.com/cs416

Suppose we trained a classifier and computed its confusion matrix on the training dataset. **Is there a class imbalance in the dataset and if so, which class has the highest representation?**

		Predicted Label		
		Pupper	Doggo	Woofers
True Label	Pupper	2	27	4
	Doggo	4	25	4
	Woofers	1	30	2

2:00

Logistic Regression

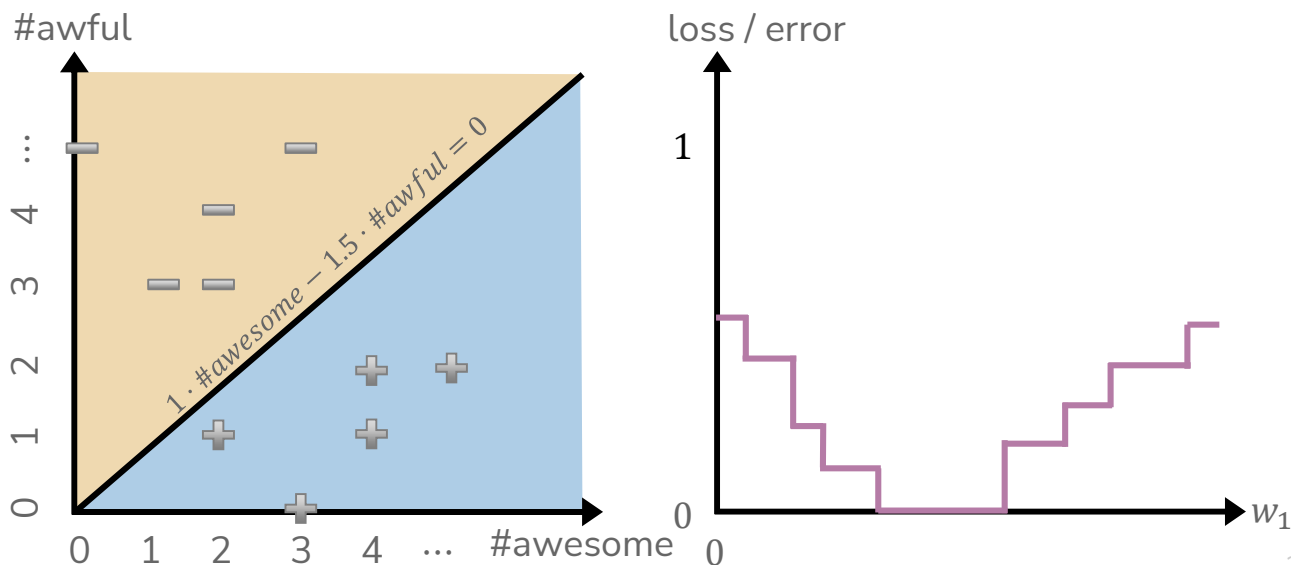
Can we use MSE for classification task?

One idea is to just model the processing of finding \hat{w} based on what we discussed in linear regression using MSE

$$\hat{w} = \underset{w}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n \mathbb{I}\{y_i \neq \hat{y}_i\}^2$$

Will this work?

Assume $h_1(x) = \#awesome$ so w_1 is its coefficient and w_2 is fixed.



Quality Metric for Classification

The MSE loss function doesn't work because of different reasons:

- The outputs are discrete values with no ordered nature, so we need a different way to frame how close a prediction is to a certain correct category
- The MSE loss function for classification task is not continuous, differentiable or convex, so we can't use optimization algorithm like Gradient Descent to find an optimal set of weights

Note: Convexity is an important concept in Machine Learning. By minimizing error, we want to find where that global minimum is, and that's ideal in a convex function.

Let's frame this problem in term of probabilities instead.

Probabilities

Assume that there is some randomness in the world, and instead will try to model the probability of a positive/negative label.

Examples:

“The sushi & everything else were awesome!”

- Definite positive (+1)
- $P(y = +1 \mid x = \text{“The sushi & everything else were awesome!”}) = 0.99$

“The sushi was alright, the service was OK”

- Not as sure
- $P(y = -1 \mid x = \text{“The sushi alright, the service was okay!”}) = 0.5$

Use probability as the measurement of certainty

$$P(y|x)$$

Probability Classifier

Idea: Estimate probabilities $\hat{P}(y|x)$ and use those for prediction

Probability Classifier

Input x : Sentence from review

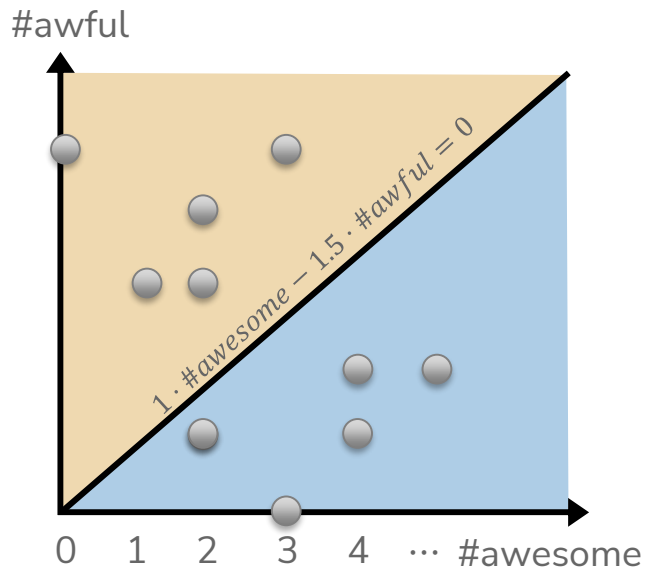
- Estimate class probability $\hat{P}(y = +1|x)$
- If $\hat{P}(y = +1|x) > 0.5$:
 - $\hat{y} = +1$
- Else:
 - $\hat{y} = -1$

Notes:

- Estimating the probability improves **interpretability**.
 - Unclear how much better a score of 5 is from a score of 3. Clear how much better a probability of 0.75 is than a probability of 0.5

Connecting Score & Probability

Idea: Let's try to relate the value of $Score(x)$ to $\hat{P}(y = +1|x)$

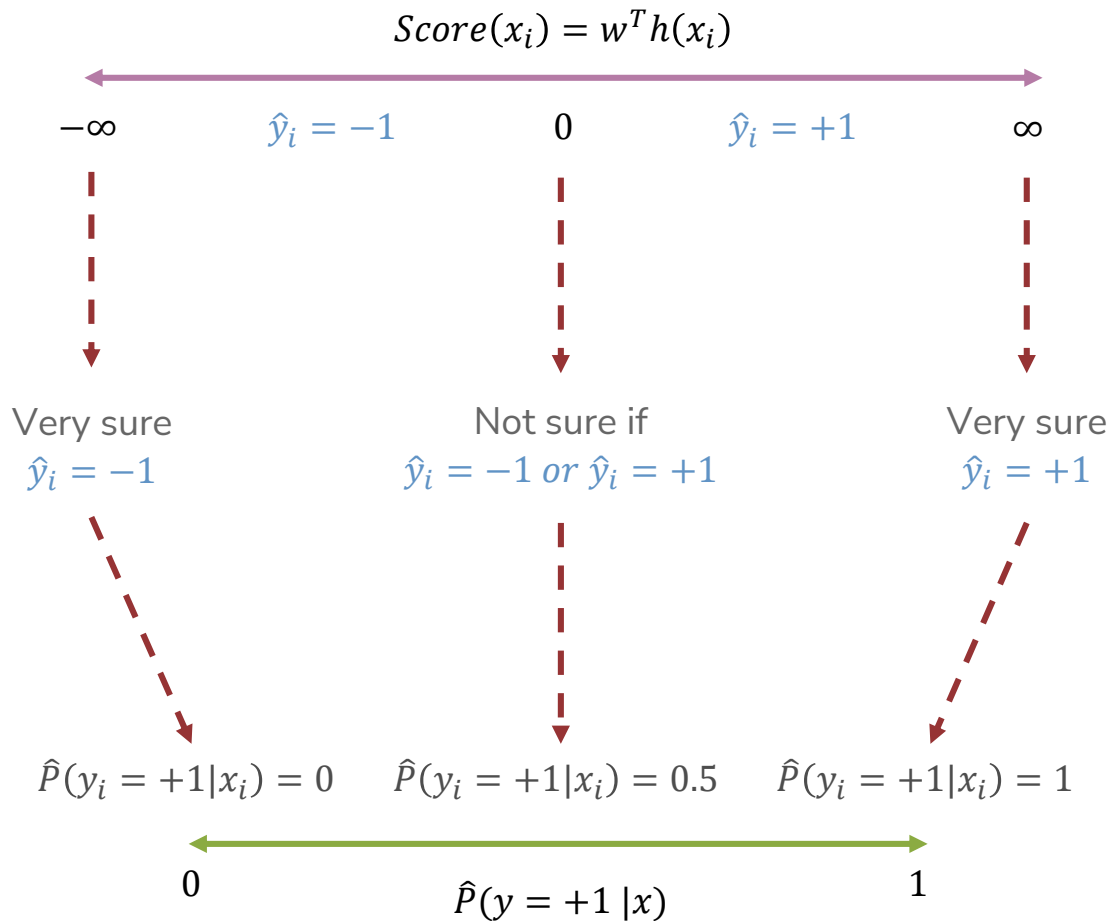


What if $Score(x)$ is positive?

What if $Score(x)$ is negative?

What if $Score(x)$ is 0?

Connecting Score & Probability

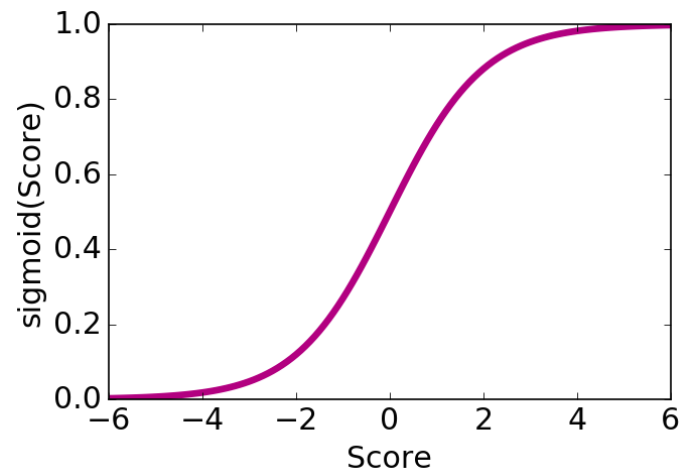


Logistic Function

Want: a function that takes numbers arbitrarily large/small and maps them between 0 and 1.

$$\text{sigmoid}(\text{Score}(x)) = \frac{1}{1 + e^{-\text{Score}(x)}}$$

$\text{Score}(x)$	$\text{sigmoid}(\text{Score}(x))$
$-\infty$	
-2	
0	
2	
∞	



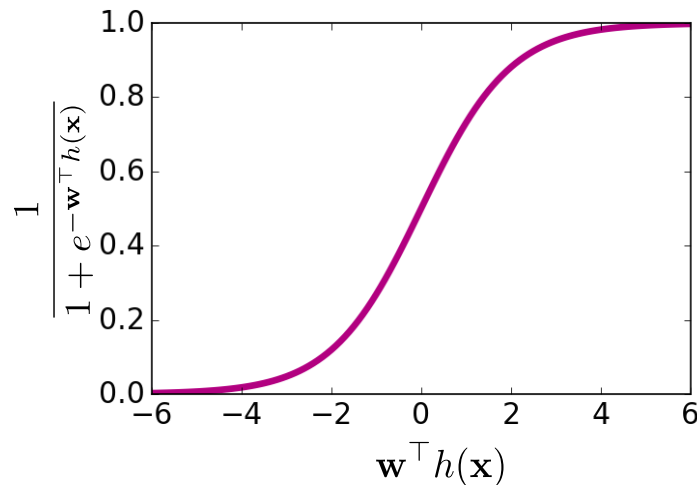
Logistic Regression Model

$$P(y_i = +1|x_i, w) = \text{sigmoid}(\text{Score}(x_i)) = \frac{1}{1 + e^{-w^T h(x_i)}}$$

Logistic Regression Classifier

Input x : Sentence from review

- Estimate class probability $\hat{P}(y = +1|x, \hat{w}) = \text{sigmoid}(\hat{w}^T h(x_i))$
- If $\hat{P}(y = +1|x, \hat{w}) > 0.5$:
 - $\hat{y} = +1$
- Else:
 - $\hat{y} = -1$



Think 

1 min

What would the Logistic Regression model predict for $P(y = -1 | x, w)$?

- "Sushi was great, the food was awesome, but the service was terrible"

- ≈ 0
- $\text{sigmoid}(-2) \approx 0.12$
- ≈ 0.5
- $\text{sigmoid}(2) \approx 0.88$
- ≈ 1

$h_1(x)$	$h_2(x)$	$h_3(x)$	$h_4(x)$	$h_5(x)$	$h_6(x)$	$h_7(x)$	$h_8(x)$	$h_9(x)$
sushi	was	great	the	food	awesome	but	service	terrible
1	3	1	2	1	1	1	1	1

Word	Weight
sushi	0
was	0
great	1
the	0
food	0
awesome	2
but	0
service	0
terrible	-1

What would the Logistic Regression model predict for $P(y = -1 | x, w)$?

- "Sushi was great, the food was awesome, but the service was terrible"
- a) ≈ 0
- b) $\text{sigmoid}(-2) \approx 0.12$
- c) ≈ 0.5
- d) $\text{sigmoid}(2) \approx 0.88$
- e) ≈ 1

$h_1(x)$	$h_2(x)$	$h_3(x)$	$h_4(x)$	$h_5(x)$	$h_6(x)$	$h_7(x)$	$h_8(x)$	$h_9(x)$
sushi	was	great	the	food	awesome	but	service	terrible
1	3	1	2	1	1	1	1	1

Word	Weight
sushi	0
was	0
great	1
the	0
food	0
awesome	2
but	0
service	0
terrible	-1

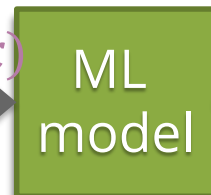
ML Pipeline



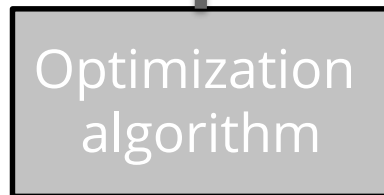
- Historical Bias
- Representation Bias
- Measurement Bias



$h(x)$



\hat{w}



$$\hat{P}(y = +1|x, \hat{w}) = \textit{sigmoid}(\hat{w}^T h(x)) = \frac{1}{1 + e^{-\hat{w}^T h(x)}}$$

Demo

Show logistic demo (see course website)





Brain Break



ML Pipeline



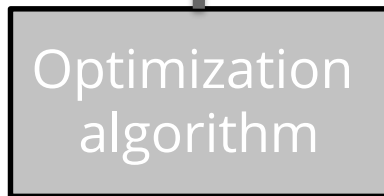
- Historical Bias
- Representation Bias
- Measurement Bias



$h(x)$



\hat{w}



$$\hat{P}(y = +1|x, \hat{w}) = \textit{sigmoid}(\hat{w}^T h(x)) = \frac{1}{1 + e^{-\hat{w}^T h(x)}}$$

Quality
Metric =
Maximum
Likelihood
Estimate

Quality Metric = Likelihood

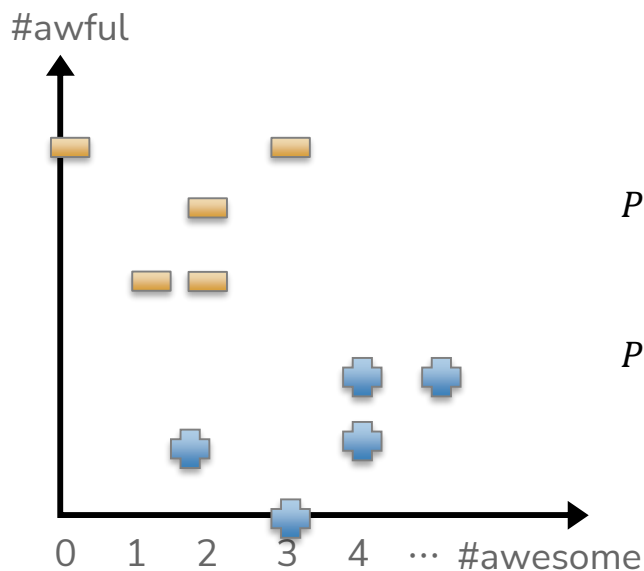
Want to compute the probability of seeing our dataset for every possible setting for w . Find w that makes data most likely!

Data Point	$h_1(x)$	$h_2(x)$	y	Choose w to maximize
$x^{(1)}, y^{(1)}$	2	1	+1	$P(y^{(1)} = +1 x^{(1)}, w)$
$x^{(2)}, y^{(2)}$	0	2	-1	$P(y^{(2)} = -1 x^{(2)}, w)$
$x^{(3)}, y^{(3)}$	3	3	-1	$P(y^{(3)} = -1 x^{(3)}, w)$
$x^{(4)}, y^{(4)}$	4	1	+1	$P(y^{(4)} = +1 x^{(4)}, w)$

Learn \hat{w}

Now that we have our new model, we will talk about how to choose \hat{w} to be the “best fit”.

- The choice of w affects how likely seeing our dataset is



$$\ell(w) = \prod_i^n P(y^{(i)}|x^{(i)}, w)$$

$$P(y^{(i)} = +1|x^{(i)}, w) = \frac{1}{1 + e^{-w^T h(x^{(i)})}}$$

$$P(y^{(i)} = -1|x^{(i)}, w) = \frac{e^{-w^T h(x^{(i)})}}{1 + e^{-w^T h(x^{(i)})}}$$

Maximum Likelihood Estimate (MLE)

Find the w that maximizes the likelihood

$$\hat{w} = \operatorname{argmax}_w \ell(w) = \operatorname{argmax}_w \prod_{i=1}^n P(y_i | x_i, w)$$

Generally, we maximize the log-likelihood which looks like

$$\hat{w} = \operatorname{argmax}_w \ell(w) = \operatorname{argmax}_w \log(\ell(w)) = \operatorname{argmax}_w \sum_{i=1}^n \log(P(y_i | x_i, w))$$

Also commonly written by separating out positive/negative terms

$$\hat{w} = \operatorname{argmax}_w \sum_{i=1: y_i=+1}^n \ln\left(\frac{1}{1 + e^{-w^T h(x)}}\right) + \sum_{i=1: y_i=-1}^n \ln\left(1 - \frac{1}{1 + e^{-w^T h(x)}}\right)$$



Likelihood vs Error/Loss

- In understanding how to measure error for the classification problem, we want to understand how close a prediction is to the correct class, which means we want to assign a high probability for a correct prediction, and low probability for an incorrect prediction
- Likelihood and error are the inverse of each other:

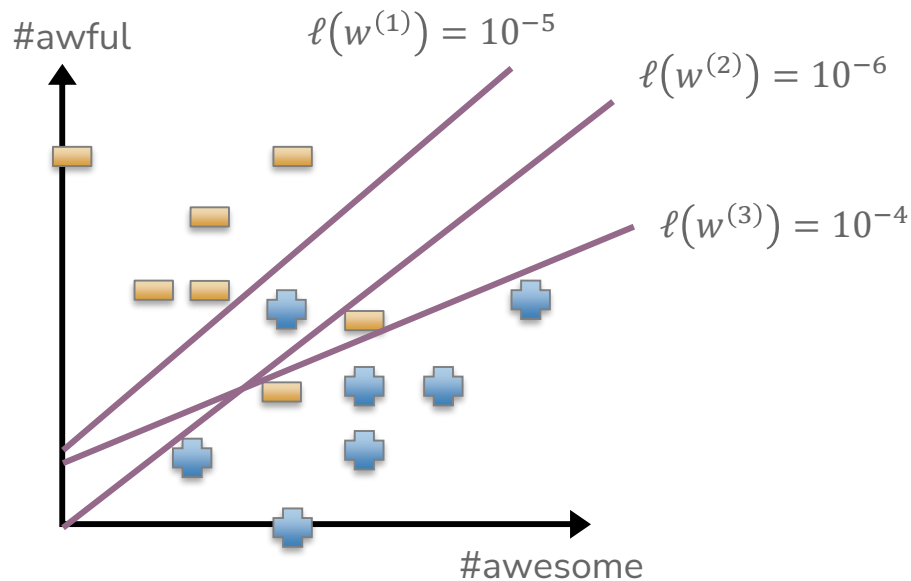
Maximizing likelihood = Minimizing Error



Think 

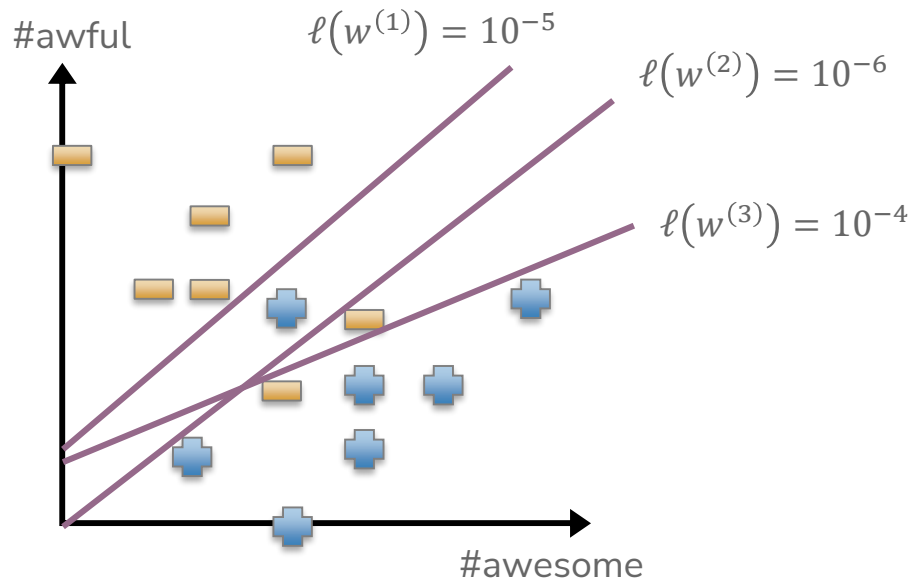
1 min

Which setting of w should we use?



pollev.com/cs416

Which setting of w should we use?



Revisiting Gradient Descent / Ascent

ML Pipeline



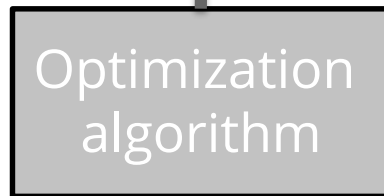
- Historical Bias
- Representation Bias
- Measurement Bias



$h(x)$



\hat{w}



$$\hat{P}(y = +1|x, \hat{w}) = \textit{sigmoid}(\hat{w}^T h(x)) = \frac{1}{1 + e^{-\hat{w}^T h(x)}}$$

Is Gradient Descent Really Used in Linear Regression?

- No!
- It **can be**, but isn't in practice.
- Linear regression has a closed form solution. The best weights are:

$$\hat{w} = (XX^T)^{-1}X^Ty$$

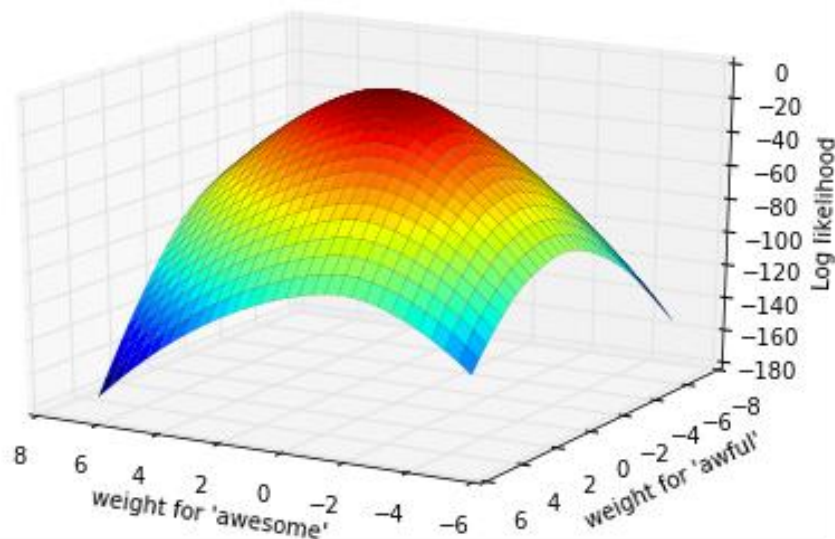
- You don't need to know the formula. What you need to know is that for Linear Regression a closed-form solution, or a solution we can write out with simple mathematical expressions, exists.
- This is not the case with Logistic Regression.
We must use Gradient Ascent/Descent!

Finding MLE

No closed-form solution, have to use an iterative method.

Since we are maximizing likelihood, we use gradient ascent.

$$\hat{w} = \operatorname{argmax}_w \prod_{i=1}^n P(y_i | x_i, w)$$



Gradient Ascent

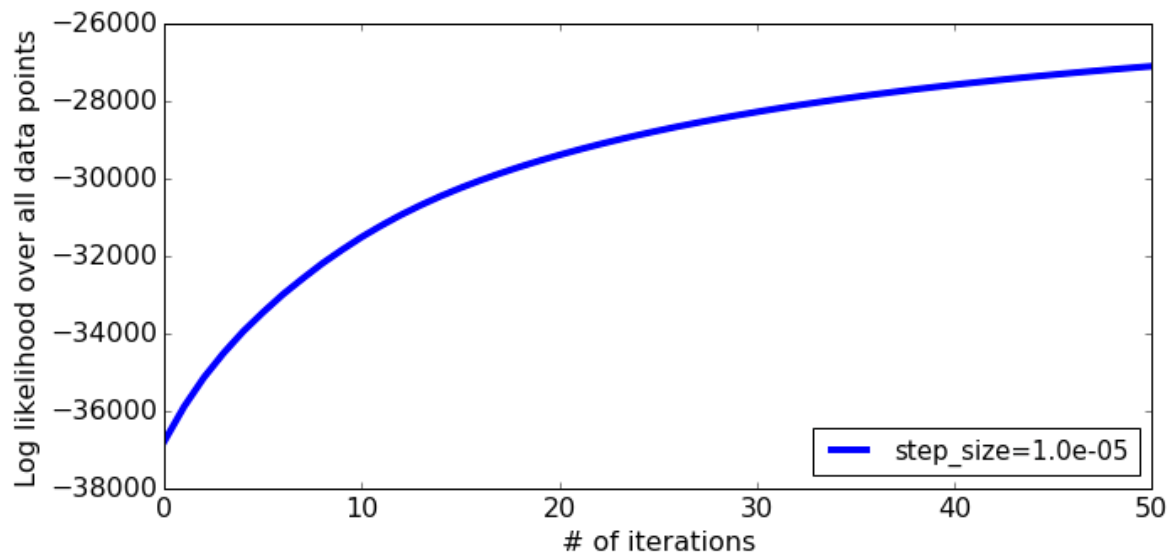
Gradient ascent is the same as gradient descent, but we go "up the hill".

```
start at some (random) point  $w^{(0)}$  when  $t = 0$   
while we haven't converged  
     $w^{(t+1)} \leftarrow w^{(t)} + \eta \nabla \ell(w^{(t)})$   
     $t \leftarrow t + 1$ 
```

This is just describing going up the hill step by step.

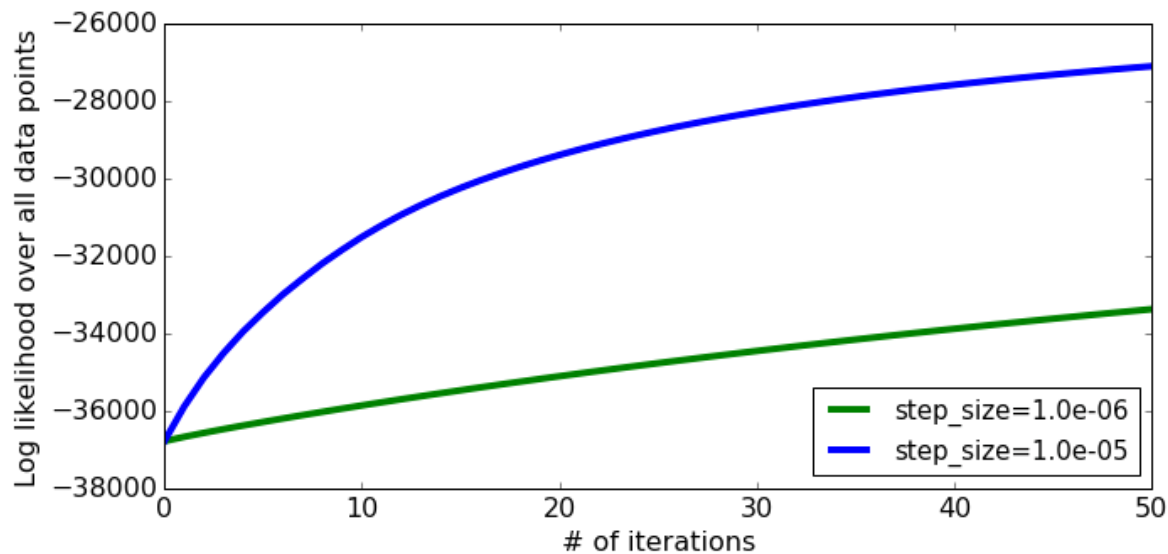
η controls how big of steps we take, and picking it is crucial for how well the model you learn does!

Learning Curve



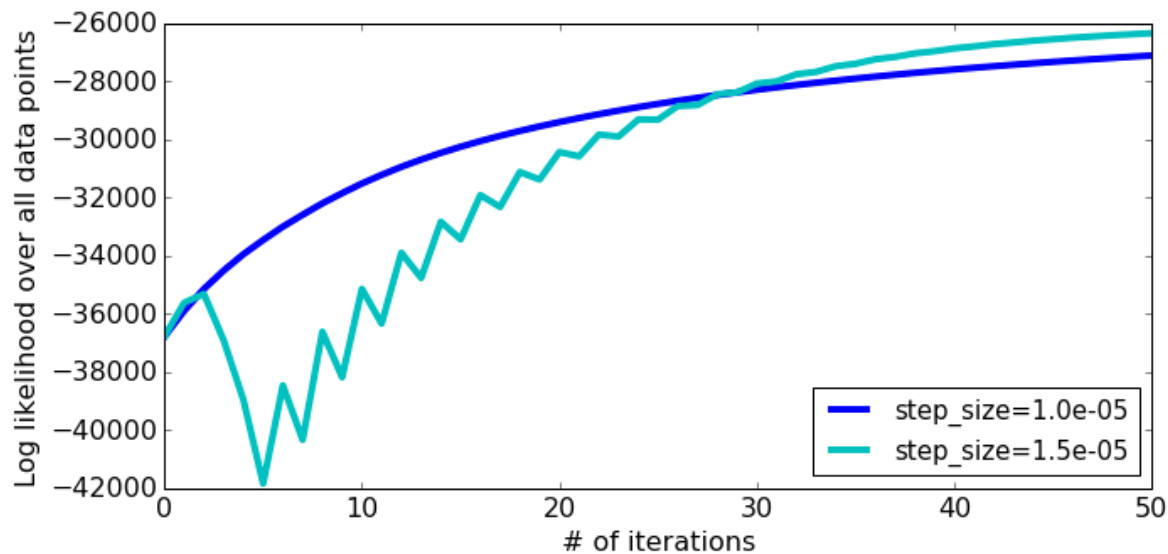
Choosing η

Step-size too small



Choosing η

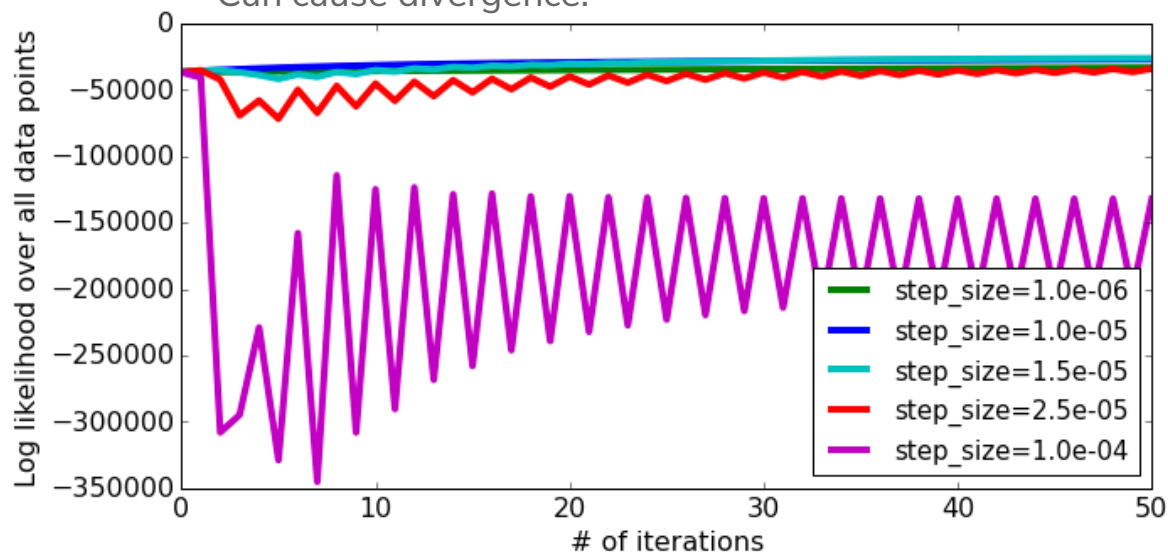
What about a larger step-size?



Choosing η

What about a larger step-size?

Can cause divergence!



Choosing η

Unfortunately, you have to do a lot of trial and error 😞

Try several values (generally exponentially spaced)

- Find one that is too small and one that is too large to narrow search range. Try values in between!

Advanced: Divergence with large step sizes tends to happen at the end, close to the optimal point. You can use a decreasing step size to avoid this

$$\eta_t = \frac{\eta_0}{t}$$



Grid Search

We have introduced yet another hyperparameter that you have to choose, that will affect which predictor is ultimately learned.

If you want to tune multiple hyperparameters at once (e.g., both a Ridge penalty and a learning rate), you will need to try all pairs of settings!

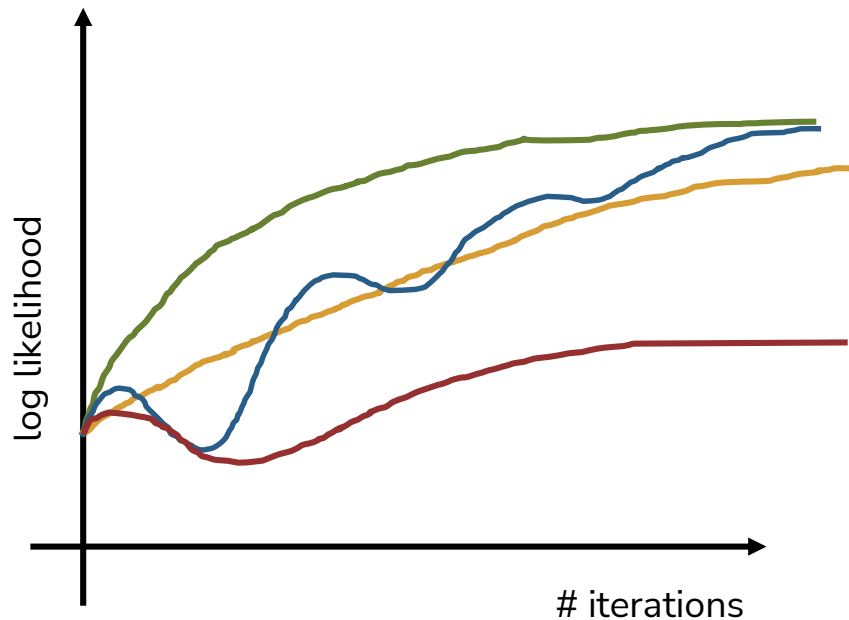
- For example, suppose you wanted to try using a validation set to select the right settings out of:
 - $\lambda \in [0.01, 0.1, 1, 10, 100]$
 - $\eta_t \in \left[0.001, 0.01, 0.1, 1, \frac{1}{t}, \frac{10}{t}\right]$
- You will need to train 30 different models and evaluate each one!



Think

1 min

- Match the below lines to the following labels:
 - “Very High Learning Rate”
 - “High Learning Rate”
 - “Good Learning Rate”
 - “Low Learning Rate”

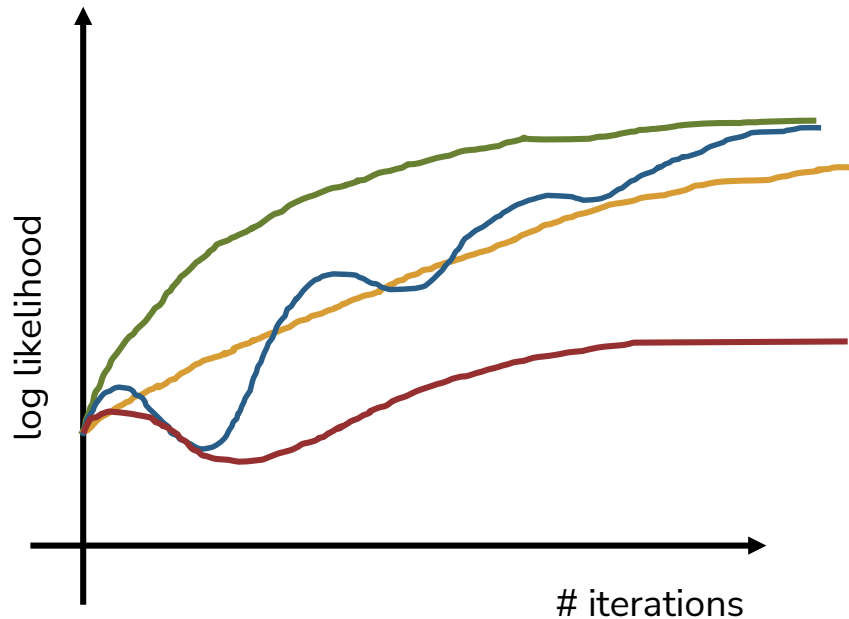


Poll Everywhere

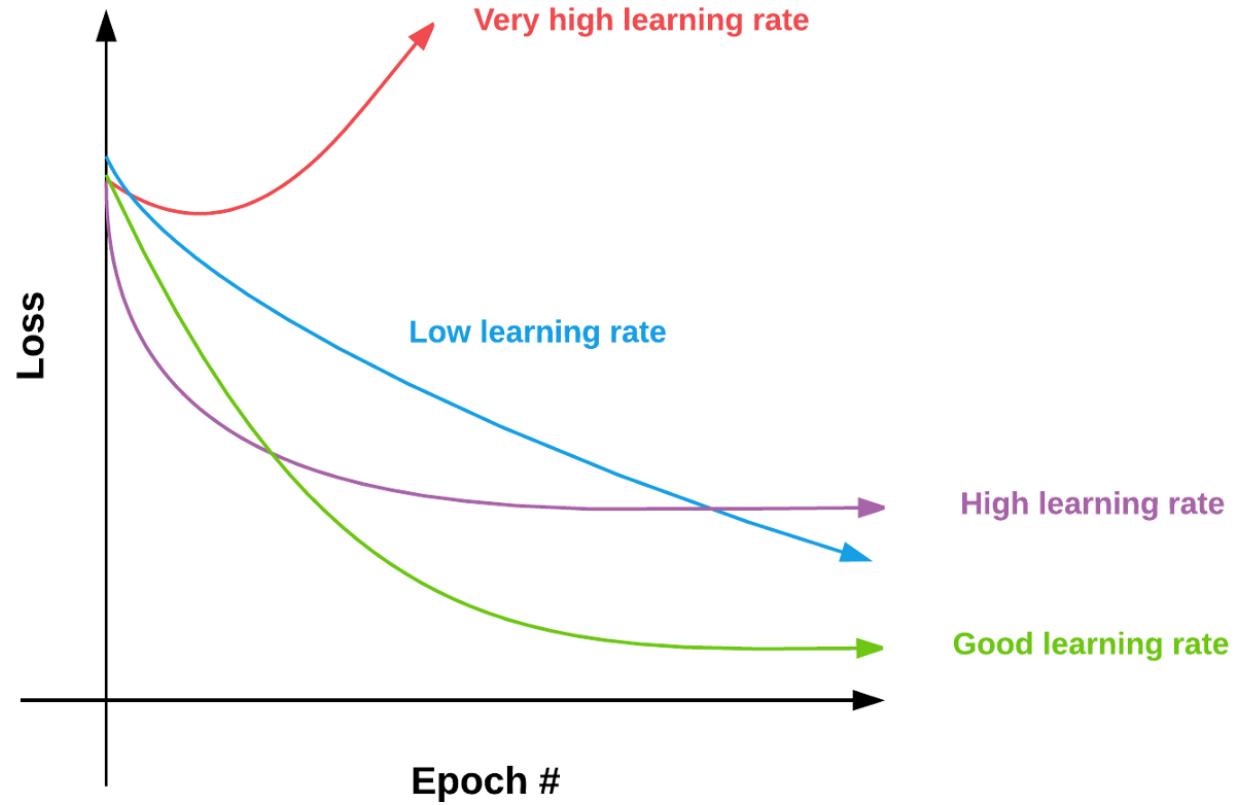
Group 

2 min

- Match the below lines to the following labels:
 - “Very High Learning Rate”
 - “High Learning Rate”
 - “Good Learning Rate”
 - “Low Learning Rate”



Likelihood vs. Loss



Overfitting - Classification

More Features

Like with regression, we can learn more complicated models by including more features or by including more complex features.

Instead of just using

$$h_1(x) = \#awesome$$

$$h_2(x) = \#awful$$

We could use

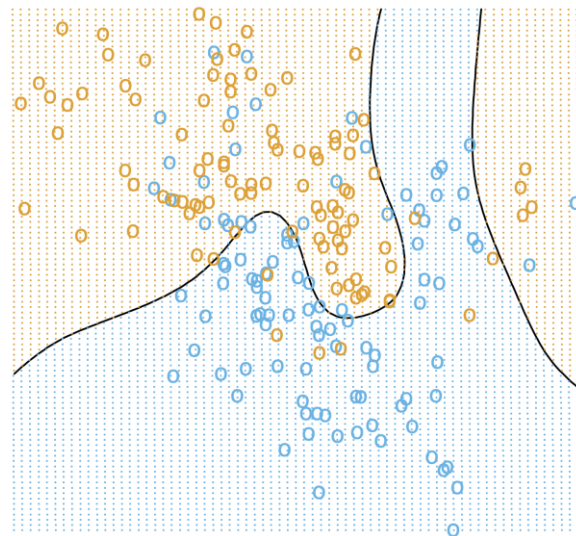
$$h_1(x) = \#awesome$$

$$h_2(x) = \#awful$$

$$h_3(x) = \#awesome^2$$

$$h_4(x) = \#awful^2$$

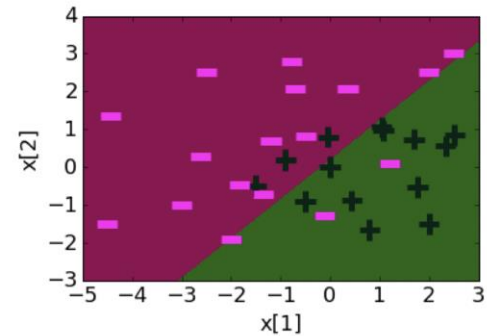
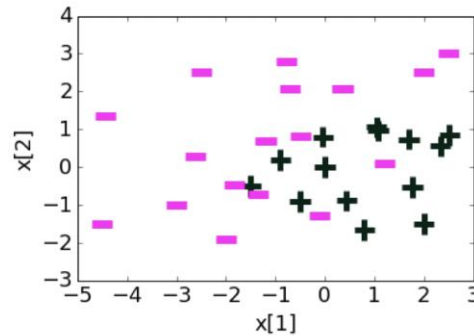
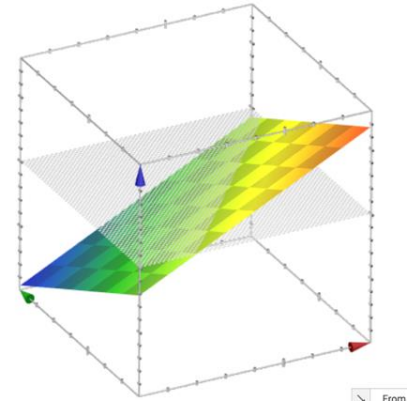
...



Decision Boundary

$$w^T h(x) = 0.23 + 1.12x[1] - 1.07x[2]$$

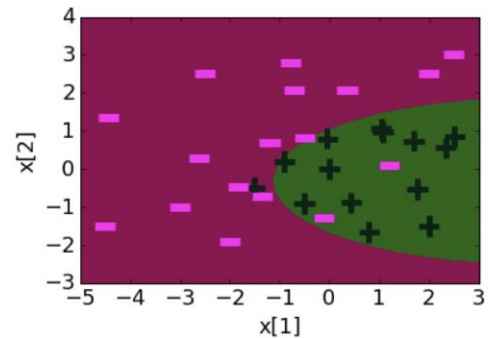
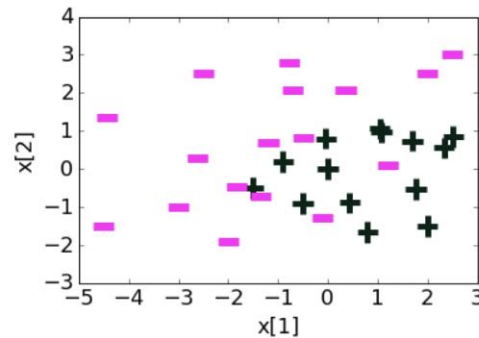
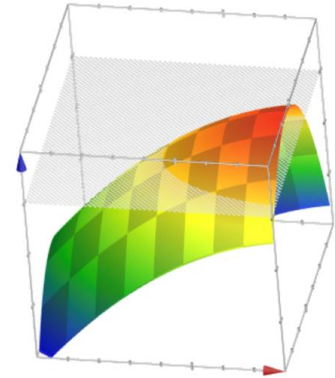
Feature	Value	Coefficient learned
$h_0(x)$	1	0.23
$h_1(x)$	$x[1]$	1.12
$h_2(x)$	$x[2]$	-1.07



Decision Boundary

$$w^T h(x) = 1.68 + 1.39x[1] - 0.59x[2] - 0.17x[1]^2 - 0.96x[2]^2$$

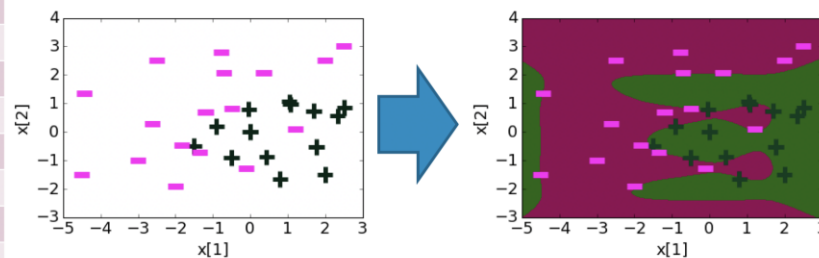
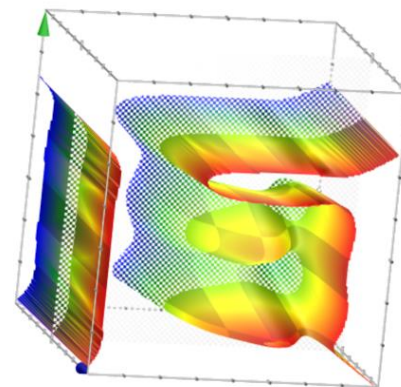
Feature	Value	Coefficient learned
$h_0(x)$	1	1.68
$h_1(x)$	$x[1]$	1.39
$h_2(x)$	$x[2]$	-0.59
$h_3(x)$	$(x[1])^2$	-0.17
$h_4(x)$	$(x[2])^2$	-0.96



Decision Boundary

$$w^T h(x) = \dots$$

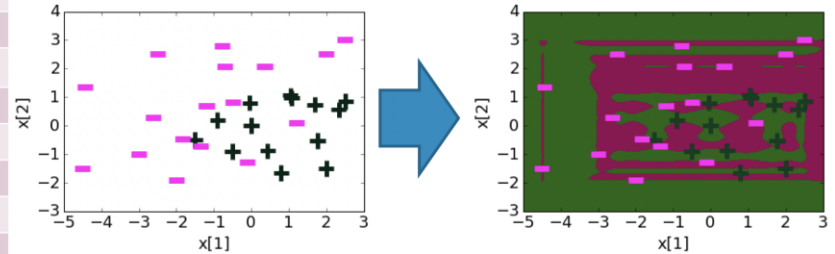
Feature	Value	Coefficient learned
$h_0(x)$	1	21.6
$h_1(x)$	$x[1]$	5.3
$h_2(x)$	$x[2]$	-42.7
$h_3(x)$	$(x[1])^2$	-15.9
$h_4(x)$	$(x[2])^2$	-48.6
$h_5(x)$	$(x[1])^3$	-11.0
$h_6(x)$	$(x[2])^3$	67.0
$h_7(x)$	$(x[1])^4$	1.5
$h_8(x)$	$(x[2])^4$	48.0
$h_9(x)$	$(x[1])^5$	4.4
$h_{10}(x)$	$(x[2])^5$	-14.2
$h_{11}(x)$	$(x[1])^6$	0.8
$h_{12}(x)$	$(x[2])^6$	-8.6



Decision Boundary

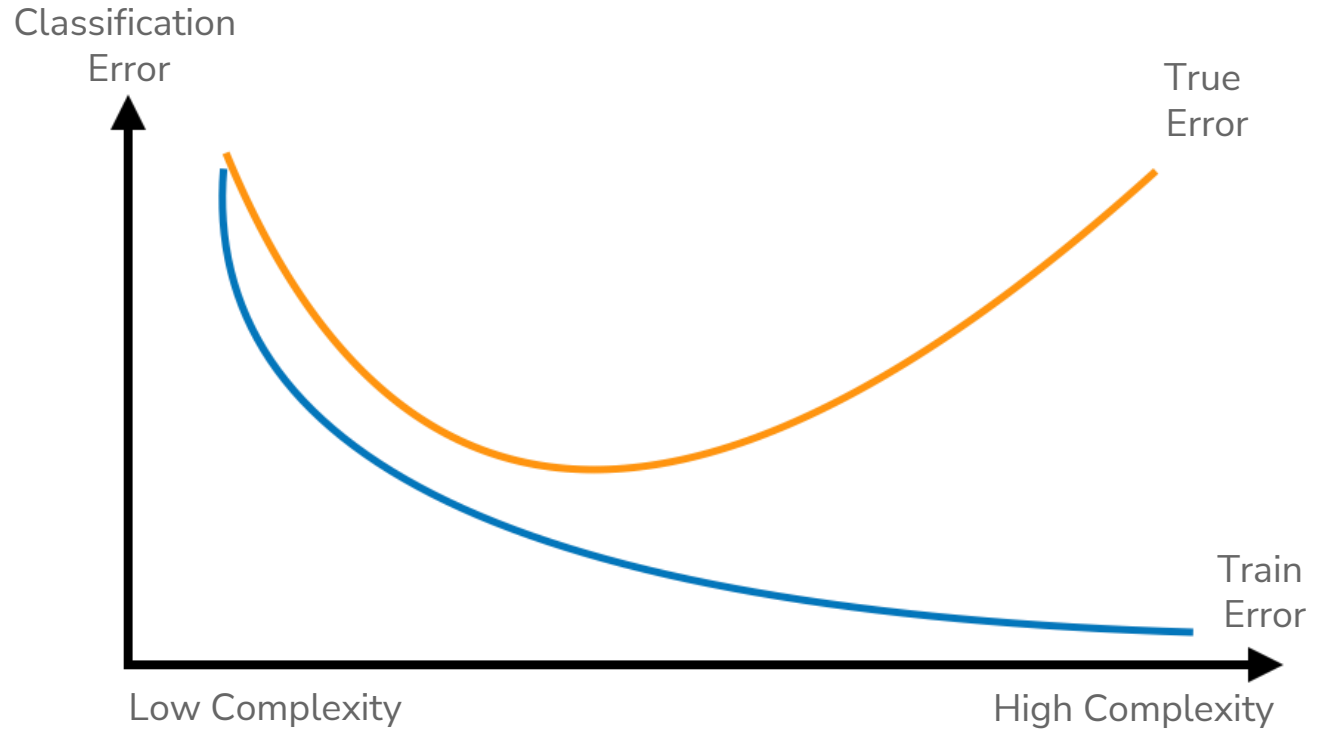
$$w^T h(x) = \dots$$

Feature	Value	Coefficient learned
$h_0(x)$	1	8.7
$h_1(x)$	$x[1]$	5.1
$h_2(x)$	$x[2]$	78.7
...
$h_{11}(x)$	$(x[1])^6$	-7.5
$h_{12}(x)$	$(x[2])^6$	3803
$h_{13}(x)$	$(x[1])^7$	21.1
$h_{14}(x)$	$(x[2])^7$	-2406
...
$h_{37}(x)$	$(x[1])^{19}$	$-2 \cdot 10^{-6}$
$h_{38}(x)$	$(x[2])^{19}$	-0.15
$h_{39}(x)$	$(x[1])^{20}$	$-2 \cdot 10^{-8}$
$h_{40}(x)$	$(x[2])^{20}$	0.03



Overfitting

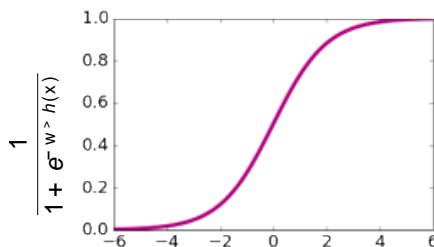
Just like with regression, we see a similar pattern with complexity



Effects of Overfitting

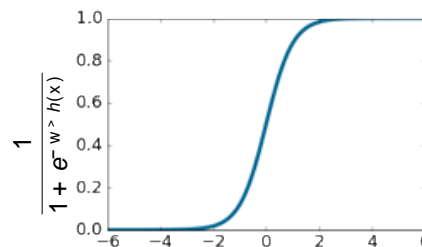
The logistic function become “sharper” with larger coefficients.

w_0	0
$w_{\text{\#awesome}}$	+1
$w_{\text{\#awful}}$	-1



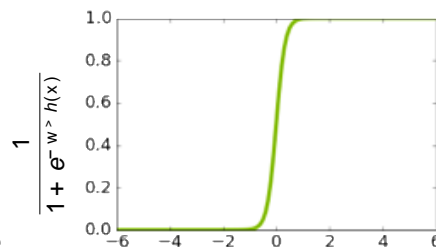
$\text{\#awesome} - \text{\#awful}$

w_0	0
$w_{\text{\#awesome}}$	+2
$w_{\text{\#awful}}$	-2



$\text{\#awesome} - \text{\#awful}$

w_0	0
$w_{\text{\#awesome}}$	+6
$w_{\text{\#awful}}$	-6



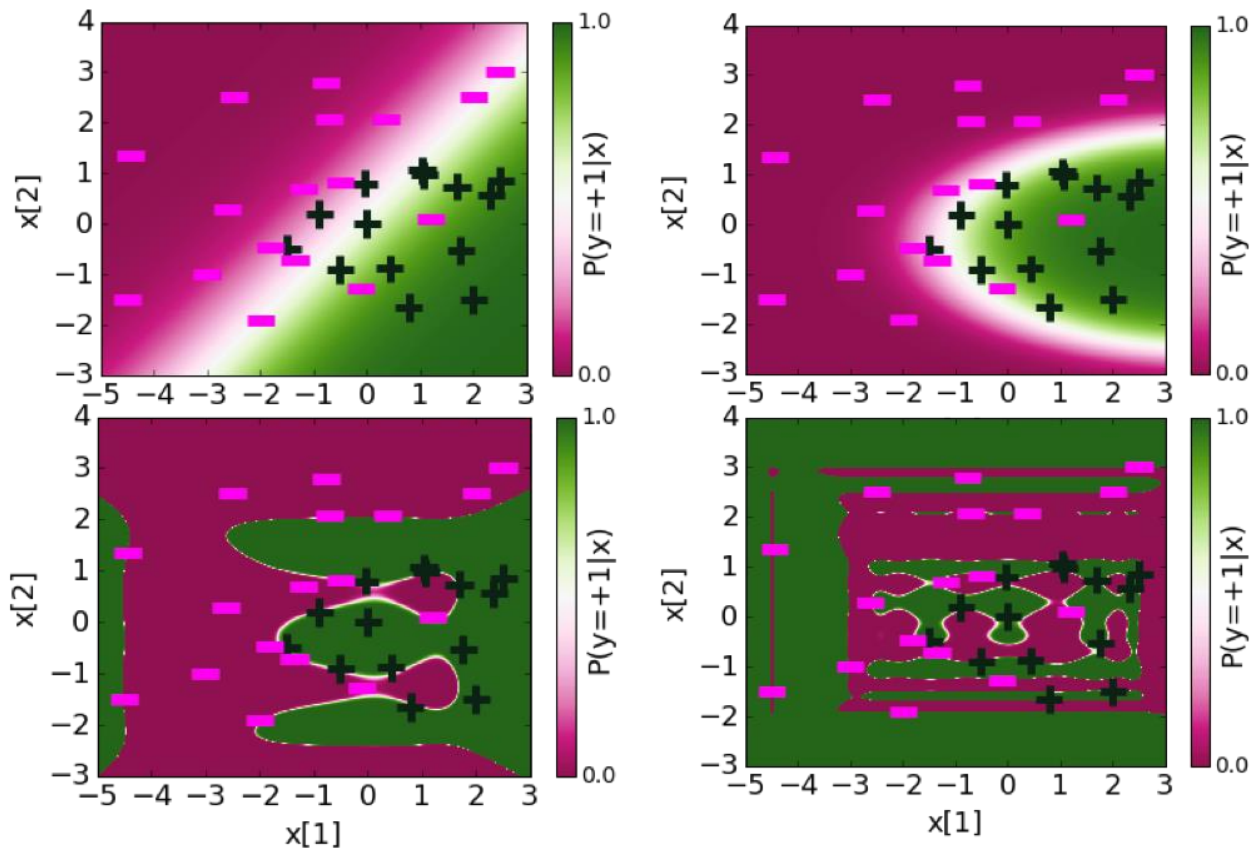
$\text{\#awesome} - \text{\#awful}$

What does this mean for our predictions?

Because the $Score(x)$ is getting larger in magnitude, the probabilities are closer to 0 or 1!

Plotting Probabilities

$$P(y = +1|x) = \frac{1}{1 + e^{-\hat{w}^T h(x)}}$$



Poll Everywhere

Think

0.5 mins

- What ideas do you have for preventing overfitting in Logistic Regression?
 - (Many possible answers)



Poll Everywhere

Group 

1.5 mins

- What ideas do you have for preventing overfitting in Logistic Regression?
 - (Many possible answers)



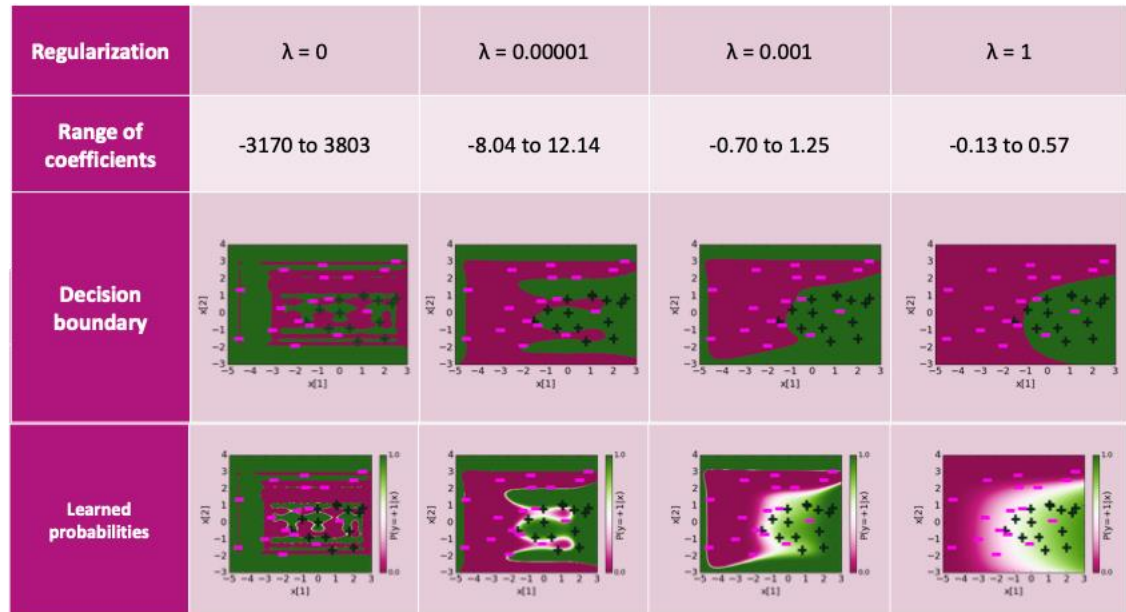
Regularization



L2 Regularized Logistic Regression

Just like in regression, can change our quality metric during training to lower the likelihood of learning an overfit model

$$\hat{w} = \operatorname{argmax}_w \ell(w) - \lambda \|w\|_2^2$$



Some Details

Why do we subtract the L2 Norm?

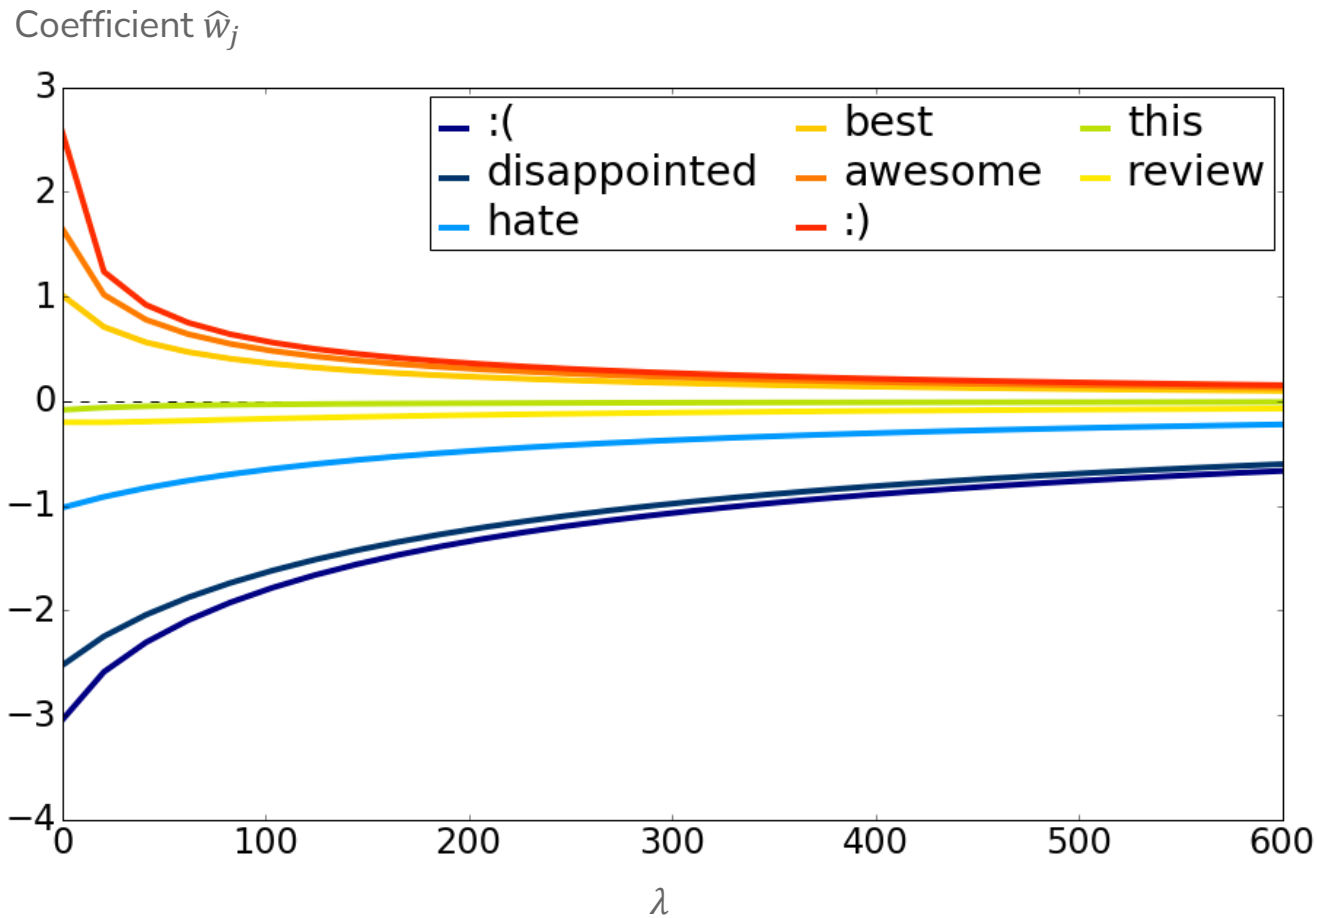
$$\hat{w} = \underset{w}{\operatorname{argmax}} \ell(w) - \lambda \|w\|_2^2$$

How does λ impact the complexity of the model?

How do we pick λ ?



Coefficient Path: L2 Penalty



Other Regularization Penalties?

Could you use the L1 penalty instead? Absolutely!

$$\hat{w} = \operatorname{argmax}_w \ell(w) - \lambda \|w\|_1$$

This is **L1 regularized logistic regression**

It has the same properties as the LASSO

- Increasing λ decreases $\|\hat{w}\|_1$
- The L1 penalty favors sparse solutions



Think

1 min

- Max wants to find the best Logistic Regression model for a sentiment analysis dataset by tuning the regularization parameter $\lambda \in [0, 10^{-2}, 10^{-1}, 1, 10]$ and the learning rate $\eta \in [10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}]$. He does the following:
 - Runs cross-validation on λ to get the best value for the regularization parameter.
 - For that value of λ , run cross-validation on η to get the best value for the learning rate.
- After running this procedure, he is convinced he has the best Logistic Regression model for his dataset, given the hyper-parameter values he wanted to test.
- **What did Max do wrong?**

Poll Everywhere

Group 

2 min

- Max wants to find the best Logistic Regression model for a sentiment analysis dataset by tuning the regularization parameter $\lambda \in [0, 10^{-2}, 10^{-1}, 1, 10]$ and the learning rate $\eta \in [10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}]$. He does the following:
 - Runs cross-validation on λ to get the best value for the regularization parameter.
 - For that value of λ , run cross-validation on η to get the best value for the learning rate.
- After running this procedure, he is convinced he has the best Logistic Regression model for his dataset, given the hyper-parameter values he wanted to test.
- **What did Max do wrong?**

Recap

Theme: Details of logistic classification and how to train it

Ideas:

- Predict with probabilities
- Using the logistic function to turn Score to probability
- Logistic Regression
- Minimizing error vs maximizing likelihood
- Gradient Ascent
- Effects of learning rate
- Overfitting with logistic regression
 - Over-confident (probabilities close to 0 or 1)
 - Regularization

